Lenovo Flex System X6 Compute Node (7903) Planning and Implementation Guide

Learn about the Flex System x880 X6, x480 X6 and x280 X6 Compute Nodes

Read our guidance about scalability and paritioning

Be informed about our latest options and features

David Watts

Simon Casey

Alexander Grechnev

LENOVO PRESS

Lenovo Flex System X6 Compute Node (7903) Planning and Implementation Guide

March 2015

SG24-8227-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Last update on March 2015

This edition applies to the Flex System X6 compute nodes, machine type 7903.

Contents

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consulty our local Lenovo representative for information on the products and services currently available in yourarea. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Advanced Settings Utility™
BladeCenter®
Bootable Media Creator™
Dynamic System Analysis™
eX5™
FlashCache™
FlashCache Storage Accelerator™
Flex System™
Lenovo®
Lenovo(logo)®
ServeRAID™
ServerGuide™
ServerProven®
System x®
ToolsCenter™
UpdateXpress System Packs™
X5™
xSeries®

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

The increasing demand for cloud computing and business analytic workloads drives innovation to find new ways to build information systems. Lenovo® clients are looking for cost-effective IT solutions that manage large amounts of data, easily scale performance, and provide reliable real-time access to actionable information.

Built on decades of innovation, Lenovo introduces its sixth generation of X-Architecture technology, the X6 servers, which includes the Flex System™ x880 X6, Flex System x480 X6, and Flex System x280 X6. These X6 servers are fast, agile, and resilient:

- *Fast* application performance means immediate access to actionable information.
- *Agile* system design helps reduce acquisition costs and provides the ability to host multiple generations of technology in a single server.
- *Resilient* platforms maximize application uptime and offer easy integration into virtual environments.

This Lenovo Press book describes the product and covers planning and implementation. The first chapters provide technical information about the 2-socket x280 X6, 4-socket x480 X6, and 8-socket x880 X6. This information is most useful in designing, configuring, and planning your server order. The later chapters provide configuration and setup details to get your server operational.

This book is for Lenovo clients, Business Partners, and employees who want to understand the features and capabilities of the X6 portfolio of servers and want to learn how to install and configure the servers for use in production. It is related to the Lenovo Press book about the rack-mount X6 servers, *System x3850 X6 and x3950 X6 Planning and Implementation Guide*, SG24-8208.

Authors

This book was produced by a team of specialists from around the world.

David Watts is a Senior IT Consultant and the program lead for Lenovo Press. He manages residencies and produces pre-sale and post-sale technical publications for hardware and software topics that are related to System x, ThinkServer, Flex System, and BladeCenter servers. He has authored over 300 books and papers. David has worked in the IT industry, both in the U.S. and Australia, since 1989, and is currently based in Morrisville, North Carolina. David holds a Bachelor of Engineering degree from the University of Queensland (Australia).

Simon Casey is an IT Specialist for the IBM Power Systems and Lenovo Flex Systems team at IBM UK, based in Hursley. With over a decade of Power Systems client experience in the financial services sector, he is now part of Lenovo's core team that implements Flex System and PureFlex System solutions, including proofs of concepts for clients. He specializes in IBM i for Flex, enterprise storage, PowerHA, and data center migrations.

Alexander Grechnev is a Lenovo Expert Technical Sales Specialist based in Moscow, Russia. He joined Lenovo in 2014 and worked for IBM for the preceding 9 years. He is responsible for presales technical support of Lenovo employees and Lenovo Business Partners and specializes in System x and Flex System products. He graduated from the Moscow State Technical University n.a. Bauman in 2002.

Thanks to the following people for their contributions to this project:

From Lenovo:

- Alfredo Aldereguia
- Ralph Begun
- Charles Clifton
- Doug Evans
- Ross Hamilton
- Kyle Hampton
- Tim Hiteshew
- Andy Huryn
- Danyel Parker
- Dave Ridley

From the IBM Redbooks organization:

- Tamikia Barrow
- Deana Coble
- Cheryl Gera
- Ilya Krutov
- Chris Rayns
- Debbie Willmschen

Portions of this book are from *System x3850 X6 and x3950 X6 Planning and Implementation Guide*, SG24-8208. Thanks to the authors:

- David Watts
- Rani Doughty
- Ilya Solovyev

1

Introduction

The Lenovo Flex System X6 family of scalable compute nodes consists of three models:

- Lenovo Flex System x880 X6 Compute Node, a 2-socket scalable compute node that scales up to an 8-socket double-wide, quad-high complex
- Lenovo Flex System x480 X6 Compute Node, a 4-socket scalable compute node that scales up to a 4-socket double-wide, double-high complex
- Lenovo Flex System x280 X6 Compute Node, an 8-socket compute node in a double-wide compute node form factor

These compute nodes are based on the sixth generation of Lenovo X-Architecture, which is based on generations of experience in designing high-end enterprise servers. The Lenovo X6 innovative design enhances scalability, reliability, availability, and serviceability features for optimal performance. They are ideal for mission-critical scalable databases, cloud applications, business analytics, virtualization, and enterprise applications.

The Flex System X6 generation compute nodes combine numerous fault-tolerant and high availability features with the Lenovo Flex System platform. This enables you to integrate multiple nodes, networking, storage, and systems management capability into a single system that is easy to deploy and manage.

The X6 Compute Nodes can be expanded on demand. This gives you flexible, modular scalability in processing, I/O, and memory so that you can provision what you need now and expand when you need to for future requirements. X6 is fast, agile, and resilient.

This chapter includes the following topics:

1.1 Target workloads

The Lenovo Flex System X6 nodes introduce new levels of fault tolerance with advanced reliability, availability, and serviceability (RAS) features implemented in hardware and software, simplified servicing and upgrades with a modular concept, dramatic improvements in response time with stretched memory speeds, and innovative flash storage offerings. These features extend the proven technologies of the previous generations of enterprise X-Architecture technology.

The X6 nodes fulfill the requirements of those who are looking for the highest level of scalable performance, maximum memory capacity, and the richest set of RAS features for maximum productivity. They are designed for mission-critical, scalable workloads, including large databases and ERP and CRM systems, to support online transaction processing (OLTP), cloud infrastructure, business analytics, virtualization, and enterprise applications.

This section describes how Lenovo X6 technology helps address challenges that Lenovo clients face in these mission-critical enterprise environments.

1.1.1 Databases

Leadership performance, scalability, and large memory support means that X6 systems can yield the best return for database applications, such as these examples:

- IBM DB2 BLU Acceleration on X6
- SAP HANA on X6
- Microsoft SQL Server Data Warehouse on X6
- SAP Business Suite on X6

X6 is well suited for OLTP workloads. OLTP workloads are characterized by small, interactive transactions that generally require sub-second response times. For most OLTP systems, the processor, memory, and I/O subsystem in a server are well balanced and not considered performance bottlenecks.

The major source of performance issues in OLTP environments is typically related to the storage I/O. The speed of traditional hard disk drive (HDD) based storage does not match the processing capabilities of the server. As a result, a situation often occurs where the processor sits idle, waiting for I/O requests to complete, which delays the user, business productivity, and infrastructure use. This is not the case with the X6.

The OLTP workload optimization of Lenovo X6 systems addresses storage I/O bottlenecks. You can choose to use system memory data, memory-channel storage, or solid-state flash-based storage.

In-memory data

Main memory (RAM) is the fastest storage type that can hold a significant amount of data. Data in main memory can be addressed by several orders of magnitude faster than data on a spinning hard disk. Using main memory as a data store is commonly known as *pinning*.

For more information about in-memory data, see 1.4, “Storage versus in-memory data” on page 6.

Solid-state storage

When sold state storage is used as a main OLTP data store, entire database structures are placed onto solid-state storage logical volumes. Sold-state storage has significantly lower latency and better IOPS performance characteristics than traditional spinning hard disk drives.

In X6, you get solid-state storage choices:

- Lenovo SSDs: Based on or 2.5-inch solid-state drives and standard SAS/SATA storage connectivity.
- External IBM FlashSystem storage systems: Based on Fibre Channel connectivity.

1.1.2 Business analytics

Data warehouses are commonly used with online analytical processing (OLAP) workloads in decision support systems, such as financial analysis. Unlike OLTP, where transactions are typically relatively simple and deal with small amounts of data, OLAP queries are more complex and process larger volumes of data.

For OLAP workloads, transactional delays can significantly increase business and financial risks. Usually, decision making is stalled or delayed because of lack of accurate, real-time operational data for analytics, which can mean missed opportunities.

These transactional delays come primarily from batch data loads and performance issues due to handling heavy complex queries and massive amounts of data (frequently referred to as *big data*) that use I/O resources. For OLAP workloads, a fast response time is critical to ensure that strategic business decisions can be made quickly in dynamic market conditions.

In general, clients might experience the following challenges with OLAP environments:

- Slow query execution and response times, which delay business decision making.
- Dramatic growth in data, which requires deeper analysis.

Lenovo X6 systems can help to make businesses more agile and analytics-driven by providing up-to-the-minute analytics based on real-time data. As with OLTP workloads, in-memory databases or flash storage is used for workload optimization.

Using Lenovo X6 technology, we help address challenges in OLAP environments in these ways:

- Dramatically boosting the performance of OLAP workloads with distributed scale-out architecture, providing almost linear and virtually unlimited performance and capacity scalability
- Significantly improving response time for better and more timely decision making

1.1.3 Virtualization

Virtualization commonly increases effectiveness in resource use, reducing capital expenses, software licensing fees, and reducing operational and management costs.

The first wave of server consolidation focused on lightly loaded servers that easily tapped into a hypervisor's ability to share processor and memory resources across applications. Hypervisors struggle to manage and share the heavy I/O loads typical of performance-intensive workloads. As a result, performance-intensive databases used for core enterprise workloads, such as enterprise resource planning (ERP), customer

relationship management (CRM), and supply chain management (SCM), are left to run on physical, non-virtualized servers.

The next wave of server virtualization with Lenovo X6 will expand the virtualization footprint to the critical applications of enterprise IT, namely those performance-intensive databases.

Lenovo X6 makes virtualization of mission-critical, performance-intensive workloads possible in several ways:

- Lenovo FlashCache™ Storage Accelerator intelligent caching software:

  This software makes it possible to virtualize high-performance databases and applications and increase per-server VM density. Deploying FlashCache Storage Accelerator™ can help you gain up to a 10x improvement in IOPS performance, up to 5x faster rich-media web page loads, reduced I/O load, and increased performance in primary storage.
- Lenovo X6 support for an integrated hypervisor:

  All X6 Compute Nodes support the addition of an internal USB key with VMware ESXi preinstalled.
- Workload optimized models:

  These are designed to take the guesswork out of deciding what the best components are for a given workload.
- Processor support:

  The Intel Xeon processor E7 v2 series supports Intel Virtualization Technology (Intel VT) Flex Priority and Intel VT Flex migration.

1.1.4 Enterprise applications: ERP and CRM

Enterprise applications, such as ERP or CRM, represent a mixed workload, where both transaction processing and certain levels of real-time reporting exist. In a 2-tier implementation, both database server and application modules are on the same server. The key performance metric is response time, akin to OLTP and OLAP workloads.

Lenovo X6 offerings provide low latency, extreme performance, and efficient transaction management to accommodate mixed workload requirements. Lenovo X6 memory-channel and flash storage can help deliver the following benefits for enterprise applications:

- Dramatically boost the performance of existing applications and lower cost-to-IOPS ratio without a need to redesign the application architecture
- Increase user productivity with better response times to improve business efficiency
- Increase data availability by using advanced system-wide high availability and reliability technologies to reduce the number of components and shorten batch processing and backup times
- Increase storage performance and capacity while decreasing power, cooling, and space requirements

1.2 Key features

The Lenovo Flex System X6 family of systems is fast, agile, and resilient and makes meeting the business needs of your enterprise easier in several ways:

- Fast application performance means immediate access to actionable information.

Lenovo X6 delivers fast application performance thanks to an innovative scalable design and new storage technology that optimizes overall performance.

- Agile system design helps reduce acquisition costs.

  Lenovo X6 delivers a unique and adaptive modular design that allows you to grow your system on demand. Scale from 2 to 4 to 8 sockets, increase your memory and I/O to meet your needs and, at the same time, reduce infrastructure cost without compromises in capacity or performance.

  With X6 technology you can gain benefits such as these:

  - You gain the ability to add I/O capability, such as extra I/O adapters or storage devices,
  - The majority of Lenovo Flex System x86 compute nodes use the same I/O adapters, so they can be interchangeable between different models of compute nodes.

- A resilient platform maximizes application uptime and promotes easy integration in virtual environments.

  The X6 Flex System architecture not only continues to operate if a component fails but also helps you reduce planned and even unplanned downtime.

  The RAS features of the new Lenovo X6 design include the following capabilities:

  - Predict failures before they happen.

    Predictive Failure Analysis allows the system to monitor the status of critical subsystems and to notify the system administrator when components appear to be degrading. Thanks to this information, in most cases, failing parts can be replaced as part of planned maintenance activity. This reduces the need for unscheduled outages, so your system continues to run.

  - Find failed components fast.

    Light path diagnostics enable systems engineers and administrators to easily and quickly diagnose hardware problems. The Chassis Management Module (CMM) and Integrated Management Module (IMM) give you even more information about the problem at hand than light path diagnostics. Failures can now be evaluated remotely, and costly downtime can be reduced or avoided altogether.

  - Survive a processor failure.

    The X6 system is designed to recover from a failed processor and restart automatically. Even if the primary node in the complex fails (the one used for booting the OS), the X6 system can boot from another processor or node by using redundant links to key resources.

  - Survive memory failures.

    The combination of Chipkill memory and redundant bit steering (RBS, also known as Double Device Data Correction, or DDDC) allows the server to tolerate two sequential DRAM memory chip failures without affecting overall system performance.

These built-in technologies drive the outstanding system availability and uninterrupted application performance that you need to host mission-critical applications.

1.3 Version comparison

The Flex System X6 Compute Nodes are part of the next generation of X-Architecture servers, as successors to the Lenovo eX5™ server line. X6 Compute Nodes offer increased scalability compared to the previous generation.

Table 1-1 on page 6 shows a high-level comparison of the 2-socket x280, 4-socket x480 X6, and 8-socket x880 X6.

Table 1-1 Large Lenovo Flex compute node comparison

Configurations		**x280 X6**	**x480 X6**	**x880 X6**
Form factor	2-socket	Double-wide (2 node bays)	Double-wide (2 node bays)	Double-wide (2 node bays)
	4-socket	Not available	Double-wide Double-high (4 node bays)	Double-wide Double-high (4 node bays)
	8-socket	Not available	Not available	Double-wide, quad-high (8 node bays)
Processors	1-node	2	2	2
	2-node	Not available	4	4
	4-node	Not available	Not available	8
Cores	1-node	30	30	30
	2-node	Not available	60	60
	4-node	Not available	Not available	120
Memory	1-node	1.5 TB	3 TB	3 TB
	2-node	Not available	6 TB	6 TB
	4-node	Not available	Not available	6 TB
2.5-inch drive bays	1-node	2	2	2
	2-node	Not available	4	4
	4-node	Not available	Not available	8
I/O adapter slots	1-node	4	4	4
	2-node	Not available	8	8
	4-node	Not available	Not available	16

1.4 Storage versus in-memory data

System main memory (RAM) is the fastest storage type that can hold a significant amount of data. Data in main memory can be accessed more than one hundred thousand times faster than data on a spinning hard disk. Even flash technology is about a thousand times slower than main memory when it comes to latency.

Main memory is connected directly to the processors through a high-speed bus, whereas hard disks or even sold-state drives are connected through a chain of buses (QPI, PCIe, SAN) and controllers (I/O hub, RAID controller or SAN adapter, and storage controller), all of which contribute to latency.

Compared to keeping data on disk, keeping data in main memory can dramatically improve database performance just by the advantage in access time. However, there is one potential

drawback: In a database transaction that has been committed, the transaction cannot stay committed.

In database design, *atomicity, consistency, isolation*, and *durability (ACID)* are a set of requirements to guarantee that database transactions are processed reliably:

- A transaction must be *atomic*. If part of a transaction fails, the entire transaction has to fail and leave the database state unchanged.
- The *consistency* of a database must be preserved by the transactions that it performs.
- *Isolation* ensures that no transaction interferes with another transaction.
- *Durability* means that after a transaction is committed, it will remain committed.

Although the first three requirements are not affected by the in-memory concept, durability is a requirement that cannot be met by storing data in main memory alone. This is because main memory is volatile storage. That is, it loses content when no electrical power is present. To make data persistent, it must either be destaged to disk (which might defeat the point of pinning to memory) or reside on non-volatile storage. Therefore, some form of permanent storage is still needed, such as hard disk drives, sold-state drives (SSDs), or flash devices to form a hybrid solution that uses both in-memory and disk technology together.

The advantage of a hybrid solution can mean flexibility by being able to balance the performance, cost, persistence and form factor, in the following ways:

- Performance: Use in-memory technology to enhance performance of sorting, storing, and retrieving specified data rather than going out to disk.
- Cost: Less costly hard disks can be substituted for more memory.
- Persistence and form factor: Memory cannot approach the density of a small hard disk.

In the next section, we discuss Lenovo technologies for Flash storage.

1.5 Energy efficiency

The X6 Compute Node design offers the following energy-efficiency features to save energy, reduce operational costs, increase energy availability, and contribute to a green environment:

- Energy-efficient electronic components help lower operational costs.
- Highly efficient 2100W AC, 2500W AC power supplies in the Lenovo Flex System chassis have earned 80 PLUS Platinum certification.
- The chassis is designed to operate in ASHRAE class A3 operating environments in temperatures up to 40° C (104° F) or altitudes up to 10,000 ft (3,000 m).
- Intel Xeon processor E7-2800/4800/8800 v2 product families offer significantly better performance over the previous generation while fitting into the same thermal design power (TDP) limits.
- Intel Intelligent Power Capability powers individual processor elements on and off, as needed, to reduce power draw.
- Low-voltage Intel Xeon processors draw less energy to satisfy the demands of power and thermally constrained data centers and telecommunication environments.
- Low-voltage 1.35 V DDR3 memory RDIMMs consume 15% less energy compared to 1.5V DDR3 RDIMMs.
- Solid-state drives (SSDs) consume as much as 80% less power than traditional HDDs.

- The node uses hexagonal ventilation holes, which is a part of Calibrated Vectored Cooling technology. Hexagonal holes can be grouped more densely than round holes, providing more efficient airflow through the system.

2

Introduction to Lenovo Flex System

The Lenovo Flex System converged infrastructure consists of hardware, software, and expertise. The Lenovo Flex System Enterprise Chassis (the major hardware component) is the next generation platform that provides new capabilities in many areas.

This chapter includes the following topics:

2.1 Lenovo Flex System Enterprise Chassis

Figure 2-1 shows the front and rear views of the Lenovo Flex System Enterprise Chassis.

Figure 2-1 Lenovo Flex System Enterprise Chassis: Front and rear

The chassis provides 14 bays for standard-width nodes, four scalable I/O switch modules, and two Chassis Management Modules (CMMs). Current node configurations include standard width and double-wide options. The chassis supports other configurations, including double-wide, double-high nodes such as the V7000 Storage Node. Power and cooling can be scaled up in a modular fashion as more nodes are added.

Table 2-1 shows the specifications of the Enterprise Chassis.

Table 2-1 Enterprise Chassis specifications

Feature	Specifications
Machine type-model	System x ordering sales channel: 8721-A1x or 8721-LRx Power Systems sales channel: 7893-92X
Form factor	10 U rack-mounted unit
Maximum number of compute nodes that are supported	14 standard (single bay), seven double-wide (two bays), or three double-height, double-wide (four bays). Intermix of node types is supported.
Chassis per 42 U rack	4
Management	One or two CMMs for basic chassis management. Two CMMs form a redundant pair; one CMM is standard in 8721-A1x. The CMM interfaces with the Integrated Management Module (IMM) or Flexible Service Processor (FSP) that are integrated in each compute node in the chassis. There is an optional IBM Flex System Manager appliance for comprehensive management, including virtualization, networking, and storage management.
I/O architecture	Up to eight lanes of I/O to an I/O adapter, with each lane capable of up to 16 Gbps bandwidth. Up to 16 lanes of I/O to a standard node with two adapters. There are a wide variety of networking solutions, including Ethernet, Fibre Channel, FCoE, RoCE, and InfiniBand.

Feature	Specifications
Power supplies	Up to six power supplies that can provide N+N or N+1 redundant power. Power supplies are 80 PLUS Platinum certified and provide over 94% efficiency at 50% load and 20% load. Each power supply contains two independently powered 40 mm cooling fan modules. Power supply options: ► 2500W AC power supply ► 2100W AC power supply ► 2500W HVDC power supply Power supplies in standard models: ► 8721-A1x: 2x 2500W AC (6 maximum) ► 8721-LRx: 2x 2100W AC (6 maximum)
Fan modules	10 fan modules (eight 80 mm fan modules and two 40 mm fan modules). Four 80 mm and two 40 mm fan modules are standard in model 8721-A1x 8721-LRx, and 7953-94X.
Dimensions	► Height: 440 mm (17.3 inches) ► Width: 447 mm (17.6 inches) ► Depth (measured from front bezel to rear of chassis): 800 mm (31.5 inches) ► Depth (measured from node latch handle to the power supply handle): 840 mm (33.1 inches)
Weight	► Minimum configuration: 96.62 kg (213 lb) ► Maximum configuration: 220.45 kg (486 lb)
Declared sound level	6.3 - 6.8 bels.
Temperature	Operating air temperature: 5 - 40° C
Electrical power	Input power: 200 - 240 V AC (nominal), 50 or 60 Hz Minimum configuration: 0.51 kVA (two power supplies) Maximum configuration: 13 kVA (six 250 W power supplies)
Power consumption	12,900 W maximum

2.2 Compute nodes

The Lenovo Flex System portfolio of servers, or *compute nodes*, includes Intel Xeon processors and IBM POWER7 and POWER7+ processors. Depending on the compute node design, the following form factors are available:

- Standard node: This node occupies one chassis bay, or half of the chassis width. An example is the Lenovo Flex System x240 Compute Node.
- Double-wide node: This node occupies two chassis bays side-by-side, the full width of the chassis. An example is the Lenovo Flex System x280 X6 Compute Node.
- Double-wide, double-high: This node occupies four chassis bays. An example is the V7000 Storage Node.

Figure 2-2 on page 12 shows a front view of the chassis, with the bay locations identified and several standard-width nodes installed.

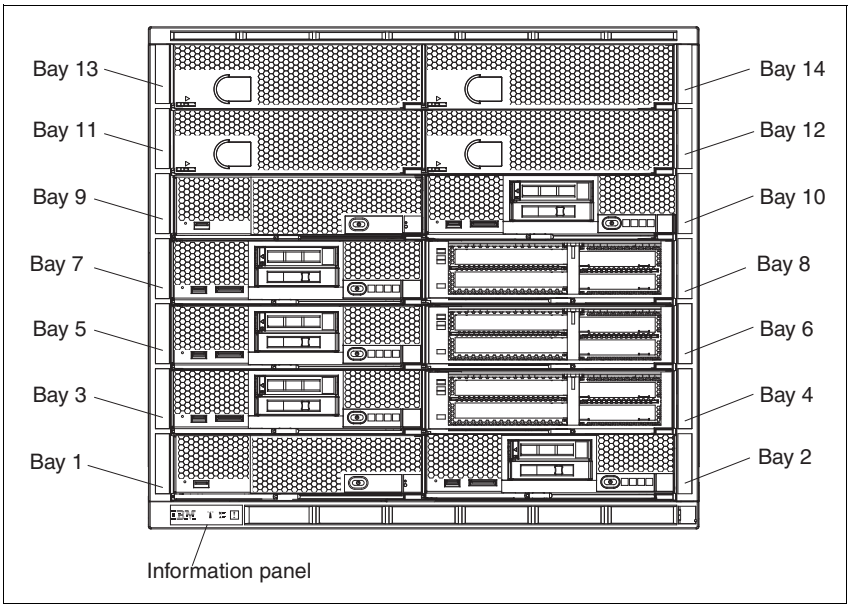

Figure 2-2 Enterprise Chassis: Front view

2.3 I/O modules

The I/O modules provide external connectivity and internal connectivity to the nodes in the chassis. These modules are scalable in terms of the number of internal and external ports that can be enabled, how these ports can be used to aggregate bandwidth, and how they can be used to create virtual switches within a physical switch. The number of internal and external physical ports that are available exceeds previous generations of products. These additional ports can be scaled or enabled as requirements grow, and more capability can be introduced.

The Enterprise Chassis can accommodate a total of four I/O modules, which are installed in a vertical orientation into the rear of the chassis, as shown in Figure 2-3 on page 13.

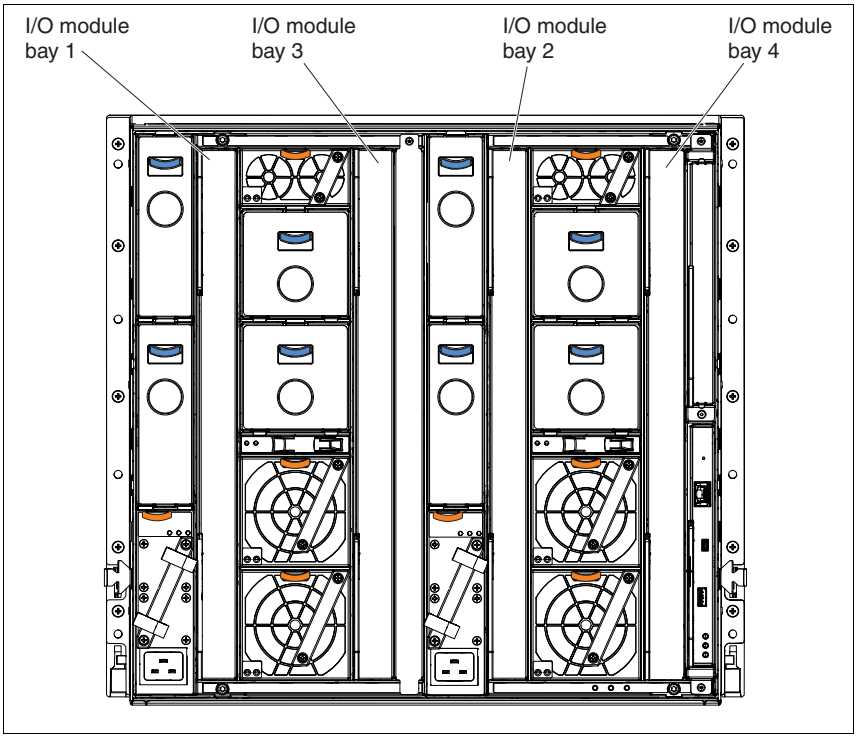

Figure 2-3 Enterprise Chassis I/O module locations

The internal connections between the node ports and the I/O module internal ports are defined by the following components:

- I/O modules 1 and 2:

  These modules connect to the ports on an I/O expansion card in slot position 1 for standard width compute nodes (such as the x240) or slot positions 1 and 3 for double-wide compute nodes (such as the x280).

  x86-based computer nodes: Most x86-based compute nodes offer integrated local area network (LAN) networking through LAN-on-Motherboard (LOM) hardware. Power Systems compute nodes have no LOM capabilities and require I/O cards for network access.

- I/O modules 3 and 4:

  These modules are connected to the ports on an I/O expansion card in slot position 2 for standard-width compute nodes or slot positions 2 and 4 for double-wide compute nodes.

  An example of I/O adapter to I/O module connectivity is shown in Figure 2-4 on page 14.

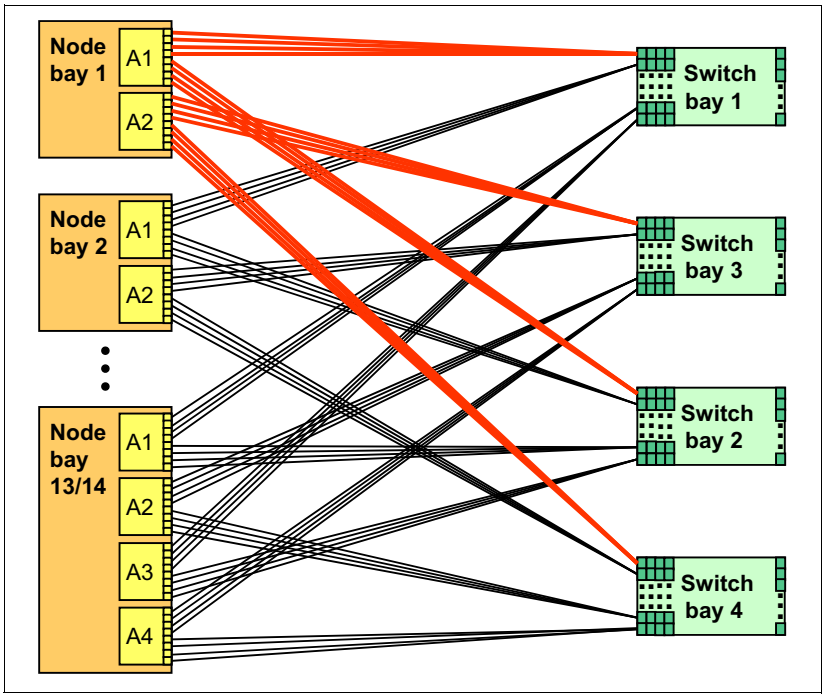

Figure 2-4 Connectivity between I/O adapter slots and switch bays

For information about available switches, see the following publications:

- Lenovo Press Product Guides for Flex System switches:

  https://lenovopress.com/servers/blades/networkmodule
- *Flex System Products and Technology*, SG24-8255:

  http://lenovopress.com/sg248255

2.4 Systems Management

Lenovo Flex System uses the following tiered approach to overall system management:

- Private management network within each chassis
- Firmware and management controllers for nodes and scalable switches
- Chassis Management Module for basic chassis management
- IBM Flex System Manager for advanced chassis management

These tiers are described next.

2.4.1 Private management network

At a physical level, the private management network is a dedicated 1 Gb Ethernet network within the chassis. This network is accessible only by the management controllers in the compute nodes or switch elements, the Chassis Management Modules, and the IBM Flex System Manager management appliance. This private network ensures a separation of the chassis management network from the data network.

The private management network is the connection for all traffic that is related to the remote presence of the nodes, delivery of firmware packages, and a direct connection to the management controller on each component.

2.4.2 Management controllers

At the next level, chassis components have their own core firmware and management controllers. Depending on the processor type of the compute nodes, an Integrated Management Module version 2 (IMMv2) or Flexible Service Processor (FSP) serves as the management controller. Additionally, each switch has a controller. In each case, the management controller provides an access point for the next level of system managers and a direct user interface.

2.4.3 Chassis Management Module

The Chassis Management Module (CMM) is a hot-swap module that is central to the management of the chassis and is required in each chassis. The CMM automatically detects any installed modules in the chassis and stores vital product data from the modules.

The CMM also acts as an aggregation point for the chassis nodes and switches, including enabling all of the management communications by Ethernet connection.

The CMM is also the key component that enables the internal management network. The CMM has a multiport, L2, 1 Gb Ethernet switch with dedicated links to all 14 node bays, the four switch bays, and the optional second CMM.

The optional second CMM provides redundancy in an active/standby mode (by using the same internal connections as the primary CMM) and is aware of all activity of the primary CMM through the trunk link between the two CMMs. This configuration ensures that the backup CMM is ready to take over in a failover situation.

2.4.4 IBM Flex System Manager

The next tier in the management stack is the IBM Flex System Manager management appliance. This is a dedicated, special-purpose, standard-width compute node that can be installed in any chassis node bay. It provides full management capabilities for up to eight chassis. All functions and software are preinstalled and are initially configured with Quick Start wizards. This integrates all components of the chassis, nodes, and I/O modules.

The Flex System Manager includes the following features:

- A single screen to manage multiple chassis and nodes
- Discovery of nodes in a managed chassis
- Integrated x86 and POWER servers, storage, and network management
- Virtualization management (VMControl)
- Upward integration to an existing IBM Tivoli environment

The Flex System Manager is a hardware appliance with a specific hardware configuration and preinstalled software stack.

Although based on a Intel compute node, the hardware platform for the Flex System Manager is not interchangeable with any other compute node. A unique expansion card that is not available on other compute nodes allows the software stack to communicate on the private management network.

The manager is available in two editions: IBM Flex System Manager and IBM Flex System Manager Advanced:

- The Flex System Manager base feature set offers the following functionality:
  - Support for up to 16 managed chassis
  - Support for up to 5,000 managed elements
  - Auto-discovery of managed elements
  - Overall health status
  - Monitoring and availability
  - Hardware management
  - Security management
  - Administration
  - Network management (Network Control)
  - Storage management (Storage Control)
  - Virtual machine lifecycle management (VMControl Express)
- The Flex System Manager Advanced feature set offers all the capabilities of the base feature set and the following features:
  - Image management (VMControl Standard)
  - Pool management (VMControl Enterprise)

2.5 Power supplies

A minimum of two and a maximum of six power supplies can be installed in the Enterprise Chassis, as shown in Figure 2-5 on page 16. All power supply modules are combined into a single power domain in the chassis that distributes power to each of the compute nodes and I/O modules through the Enterprise Chassis midplane.

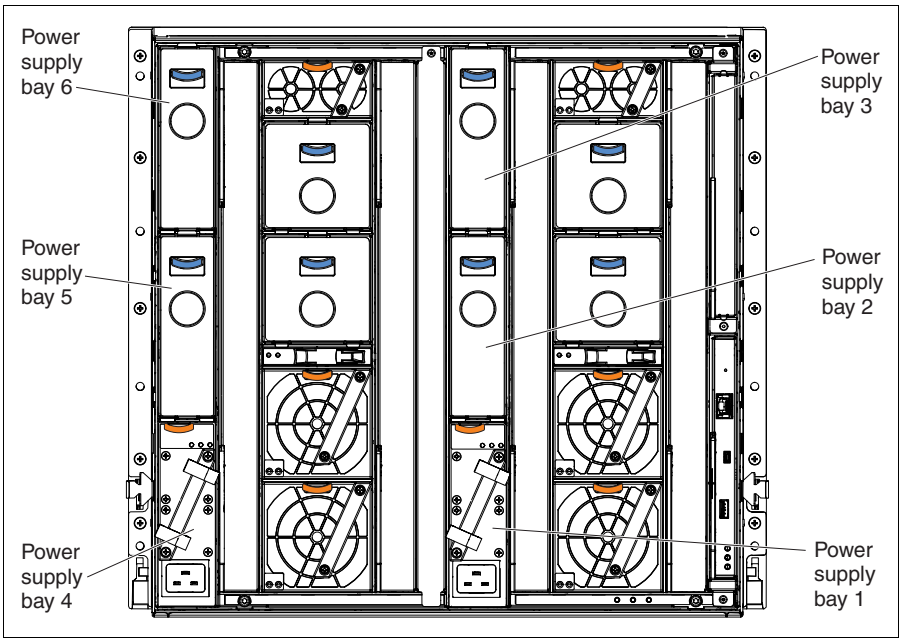

Figure 2-5 Enterprise Chassis power supply locations

Currently, the following types of power supplies are available:

- 2100 W AC power supplies
- 2500 W AC power supplies
- 2500 W HVDC power supplies

Table 2-2 shows the maximum number of configurable X6 Compute Nodes for the power supplies that are installed in the chassis. The following color codes are used in the table:

- Green: No restriction to the number of compute nodes installable
- Yellow: Some restrictions apply and some bays must be left unpopulated

Table 2-2 Specific number of compute nodes supported based on installed power supplies

Compute node	CPU TDP rating	2100W power supplies (AC)				2500W power supplies (AC or DC)			
		N+1, N=5 6 total	N+1, N=4 5 total	N+1, N=3 4 total	N+N, N=3 6 total	N+1, N=5 6 total	N+1, N=4 5 total	N+1, N=3 4 total	N+N, N=3 6 total
x280, x480 or x880 X6 with 2 sockets	105 W	7	7	6	6	7	7	7	7
	130 W	7	7	6	6	7	7	7	7
	155 W	7	7	5	6	7	7	7	7
x480 or x880 X6 with 4 sockets	105 W	3	3	3	3	3	3	3	3
	130 W	3	3	3	3	3	3	3	3
	155 W	3	3	2	3	3	3	3	3
x880 X6 with 8 sockets	105 W	1	1	1	1	1	1	1	1
	130 W	1	1	1	1	1	1	1	1
	155 W	1	1	1	1	1	1	1	1

Power configurator: For more information about exact configuration, see the Lenovo System x Power Configurator:

`http://www.ibm.com/systems/bladecenter/resources/powerconfig.html`

The 2100 W AC and 2500 W AC power supplies are 80 PLUS Platinum certified.

The *80 PLUS Platinum* standard is a performance specification for power supplies that are used in servers and computers. To meet this standard, the power supply must have an energy efficiency rating of 90% or greater at 20% of rated load, 94% or greater at 50% of rated load, and 91% or greater at 100% of rated load, with a power factor of 0.9 or greater. For more information about the 80 PLUS Platinum standard, see this website:

`http://www.80PLUS.org`

The Enterprise Chassis allows configurations of power policies to give N+N or N+1 redundancy.

Tip: N+1 in this context means a single backup device for N number of devices. Any component can replace any other component, but only once.

N+N means that there are N backup devices for N devices, where N number of devices can fail and each has a backup.

2.6 Power policies

The redundancy options are configured from the CMM and can be changed nondisruptively. The five policies are shown in Table 2-3.

Table 2-3 Chassis power management policies

Power management policy	Function
Basic Power Management	Allows the chassis to fully use available power (no N+N or N+1 redundancy).
Power Module Redundancy	Single power supply redundancy with no compute node throttling (N+1 redundancy).
Power Module Redundancy with Compute Node Throttling allowed	Single power supply redundancy. Compute nodes can be throttled (if required) to stay within the available power. This setting provides higher power availability over simple Power Module Redundancy (N+1 setting).
Power Source Redundancy	Maximum power available, limited to one-half of the installed number of power supplies (N+N setting).
Power Source Redundancy with Compute Node Throttling allowed	Maximum power available, limited to one-half of the installed number of power supplies. Compute nodes can be throttled (if required) to stay within available power. This setting provides higher power availability compared with simple Power Source Redundancy (N+N setting).

Figure 2-6 shows the available power management policies in the CMM.

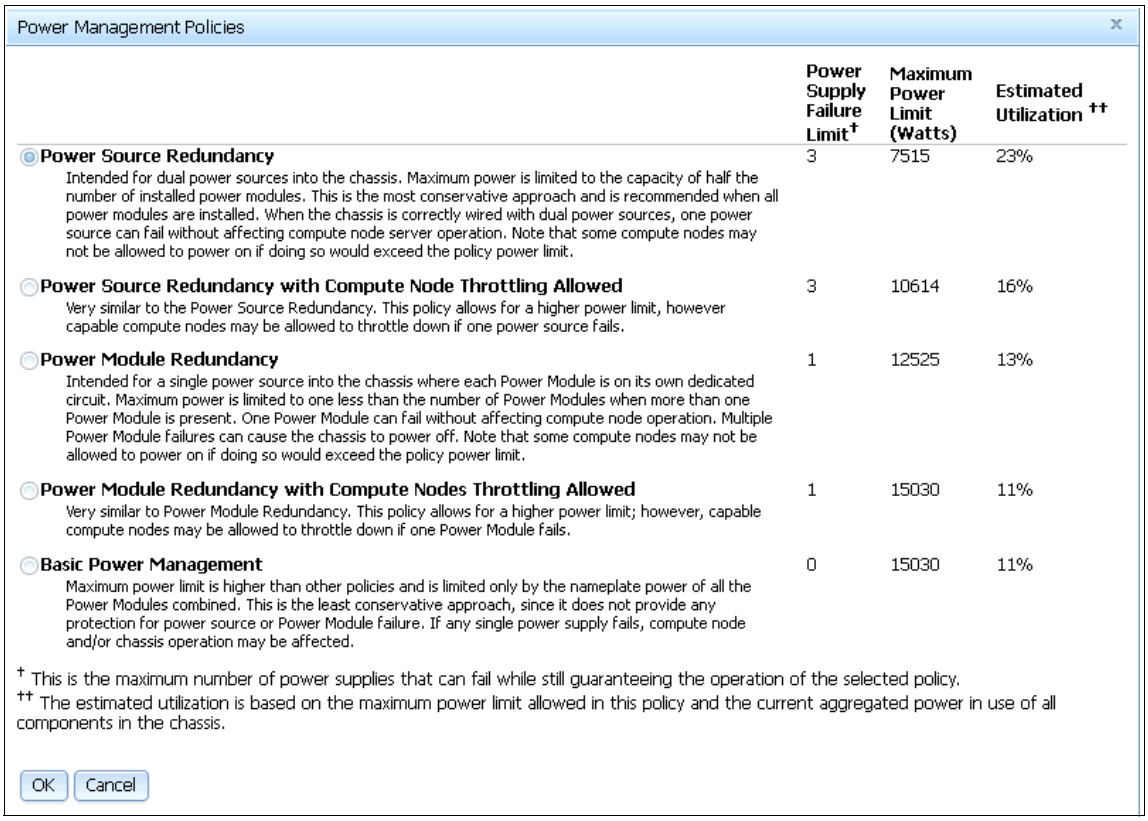

Figure 2-6 Power management policies in CMM

In addition to the redundancy settings, a power limiting and capping policy can be enabled by the CMM to limit the total amount of power that a chassis requires.

For more information about power supplies, see *Flex System Products and Technology*, SG24-8255:

http://lenovopress.com/sg248255

2.7 Cooling

The flow of air in the Enterprise Chassis follows a front to back cooling path, where cool air is drawn in at the front of the chassis and warm air is exhausted to the rear. Air movement is controlled by hot-swappable fan modules in the rear of the chassis and a series of internal dampers.

The cooling is scaled up as required, based on the number of nodes installed. (The number of cooling fan modules that is required for a number of nodes is shown in Table 2-4 on page 21.)

Chassis cooling is adaptive and node-based rather than chassis-based. Inputs into the cooling algorithm are determined from the following factors:

- Node configurations
- Power monitor circuits
- Component temperatures
- Ambient temperature

With these inputs, each fan module has greater independent granularity in fan speed control. This results in lower airflow volume (CFM) and lower cooling energy that is spent at the chassis level for any configuration and workload.

Figure 2-7 shows the location of the fan modules.

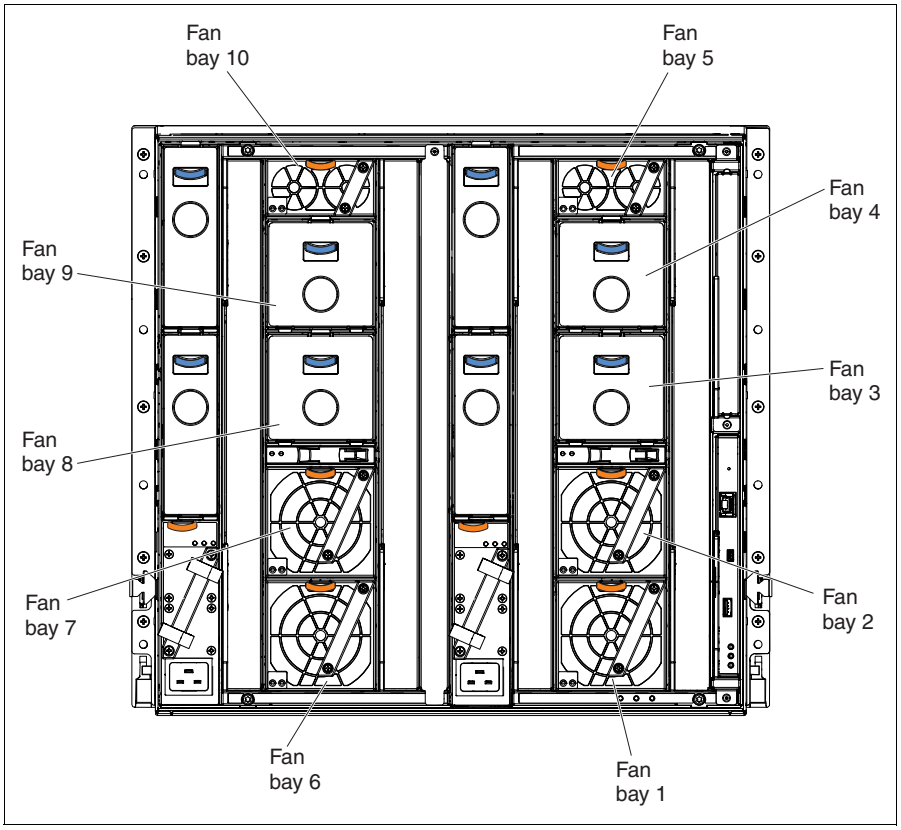

Figure 2-7 Enterprise Chassis fan module locations

2.7.1 Node cooling

There are two compute node cooling zones: zone 1 on the right side of the chassis and zone 2 on the left side of the chassis (both viewed from the rear). The chassis can contain up to eight 80 mm fan modules across the two zones. Four 80 mm fan modules are included in the base configuration for node cooling. Other fan modules are added in pairs across the two zones.

Figure 2-8 shows the node cooling zones and fan module locations.

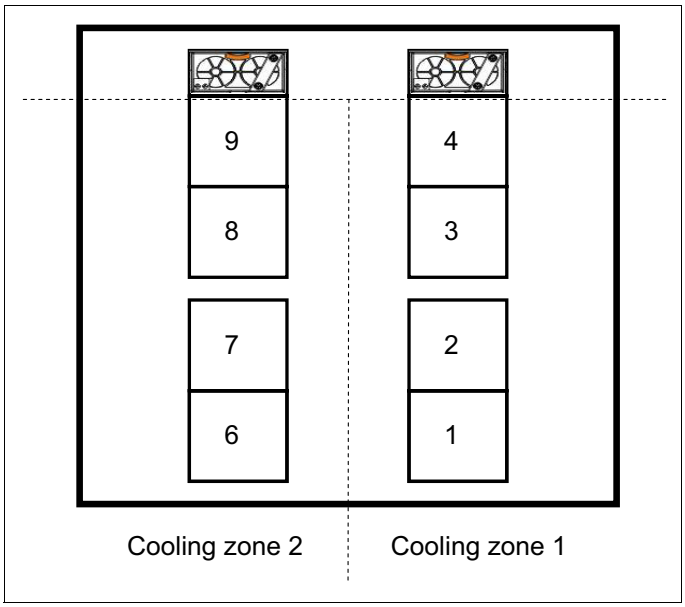

Figure 2-8 Enterprise Chassis node cooling zones and fan module locations

When a node is not inserted in a bay, an airflow damper closes in the midplane to prevent air from being drawn through the unpopulated bay. By inserting a node into a bay, the damper is opened, thus allowing cooling of the node in that bay.

Table 2-4 shows the relationship between the number of fan modules and the number of nodes supported.

Table 2-4 Fan module options and numbers of supported nodes

Fan module option	Total number of fan modules	Total number of nodes supported
Base	4	4
First option	6	8
Second option	8	14

Chassis area: The chassis node-bay area is effectively one large chamber. Nodes can be placed in any slot; however, preferred practices indicate that the nodes must be placed as close together as possible to be inline with the fan modules.

2.7.2 Switch and Chassis Management Module cooling

There are two other cooling zones for the I/O switch bays. These zones, zones 3 and 4, are at the upper right and upper left side of the bays, as viewed from the rear of the chassis. Cooling zones 3 and 4 are serviced by 40 mm fan modules that are included in the base configuration and cool the four available I/O switch bays.

Upon hot-swap removal of a 40 mm fan module, a back flow damper in the fan bay closes. The backflow damper prevents hot air from entering the system from the rear of the chassis. When the fan module is being replaced, the 80 mm fan modules cool the I/O modules and the Chassis Management Module. Figure 2-9 shows cooling zones 3 and 4 that service the I/O modules.

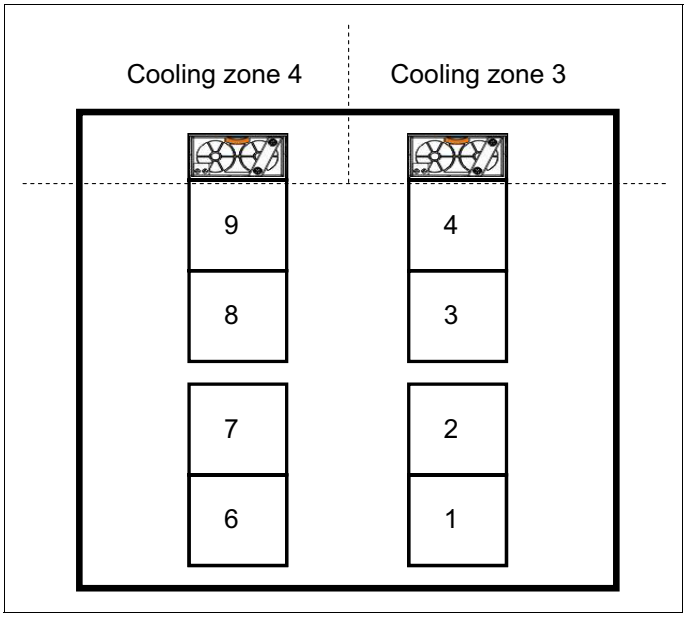

Figure 2-9 Cooling zones 3 and 4

2.7.3 Power supply cooling

The power supply modules have two integrated 40 mm fans. Installation or replacement of a power supply and fans is done as a single unit.

The integral power supply fans are not dependent upon the power supply being functional. Rather, they are powered independently from the midplane.

3

Architecture

In this chapter, we give an overview of the technologies that Lenovo has included in the Flex System X6 Compute Node family. We cover the following topics:

3.1 Modular design

The Lenovo Flex X6 Compute Node family consists of the new flagship nodes of the Lenovo Flex System x86-based family:

- Lenovo Flex System x880 X6 Compute Node, a 2-socket scalable compute node that scales:
  - Two x880 X6 servers connected to form a 4-socket double-high complex
  - Four x880 X6 servers connected to form an 8-socket quad-high complex
- Lenovo Flex System x480 X6 Compute Node, a 2-socket scalable compute node that scales:
  - Two x480 X6 servers connected to form a 4-socket double-high complex
- Lenovo Flex System x280 X6 Compute Node, a 2-socket compute node in a double-wide compute node form factor

Figure 3-1 shows two x880 X6 Compute Nodes connected to form a 4-socket server.

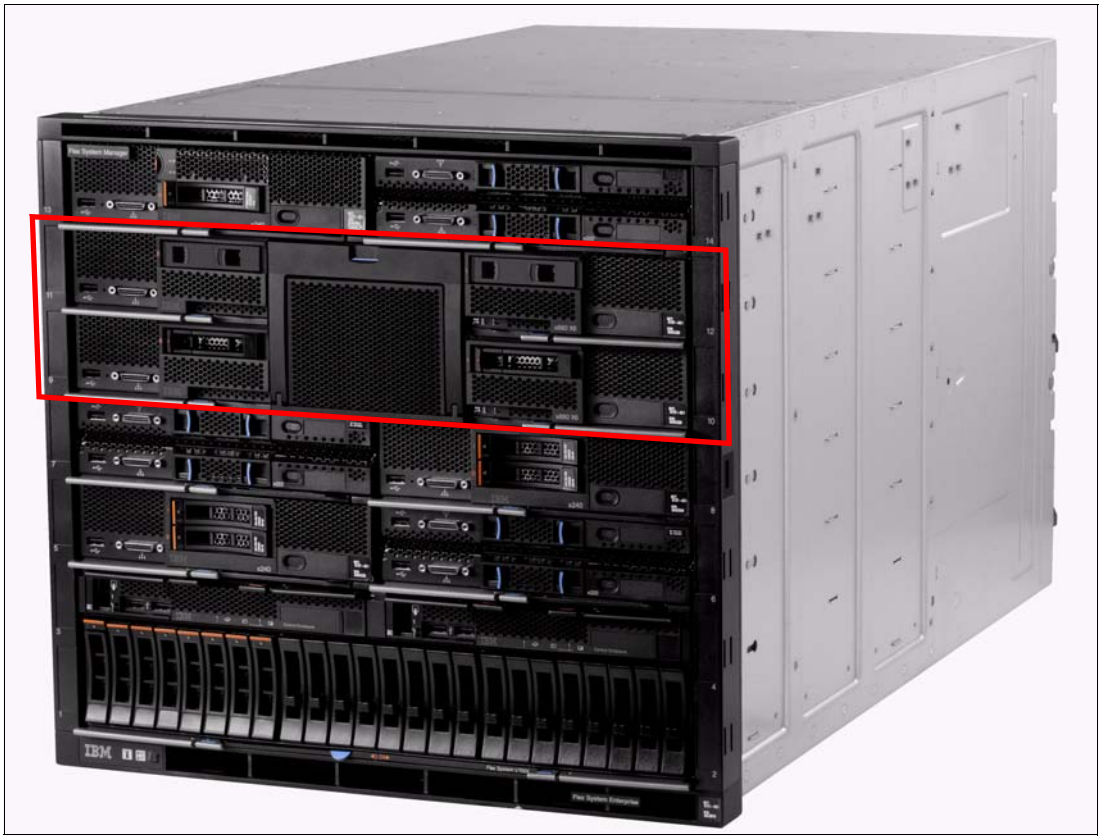

Figure 3-1 Two Flex System x880 Compute Nodes connected to form a 4-socket complex

The x280 X6, x480 X6, and x880 X6 systems all share the same fundamental building block in the node design, with the differentiator being the model of Intel Xeon processor installed:

- The x880 X6 is based on the Intel Xeon Processor E7-8800 v2 family
- The x480 X6 is based on the Intel Xeon Processor E7-4800 v2 family
- The x280 X6 is based on the Intel Xeon Processor E7-2800 v2 family

The Flex X6 Compute Node is a dual-wide 2-socket compute node. The node's width allows for significant I/O capability, providing 48 DIMM sockets for memory or storage and four I/O adapter slots for a high level of bandwidth and resiliency.

Figure 3-2 shows a Flex X6 Compute Node with the top cover removed.

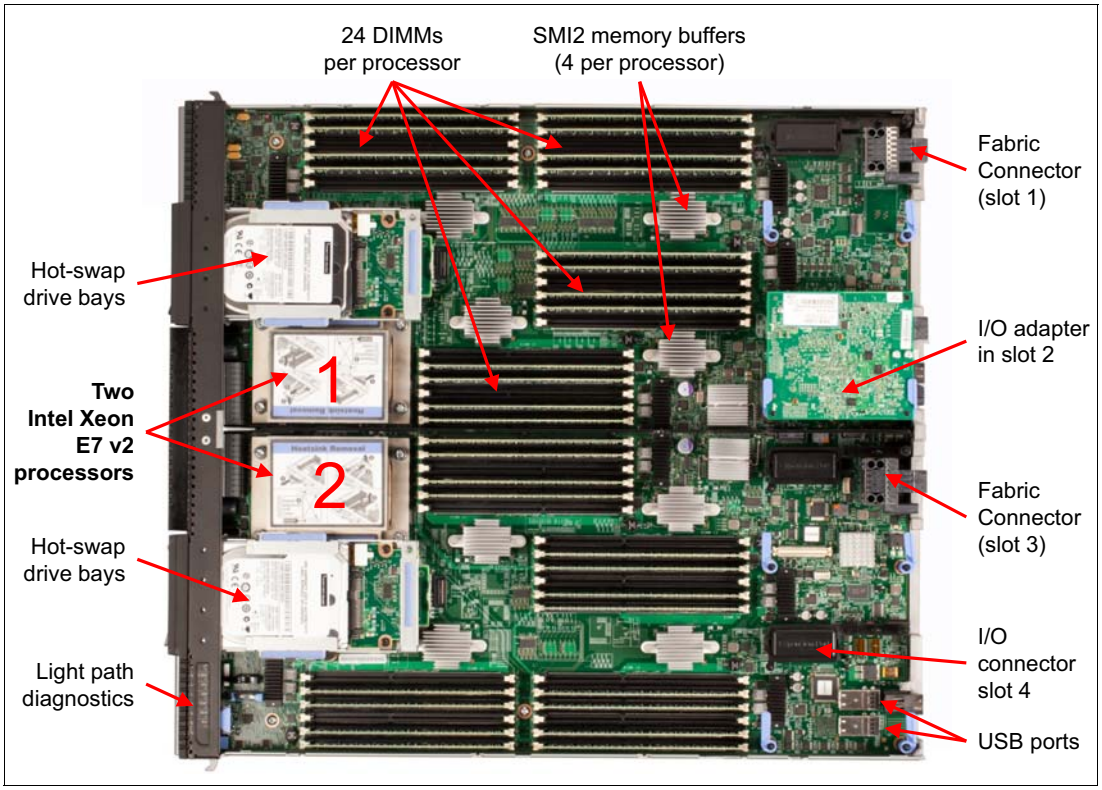

Figure 3-2 The Flex X6 Compute Node

Leveraging this common architecture allows the Flex X6 Compute Node family to be physically scaled from a 2-socket to a 4- or 8-socket complex, if the processor supports it.

3.2 Systems architecture

This section shows the overall architecture of the x280 X6, x480 X6, and x880 X6. It also covers the fundamentals of the 2-socket building block that comprise all systems.

Figure 3-3 shows the system architecture of a 2-socket X6 Compute Node. The architecture includes two QPI ports that are presented externally to the server to allow connections to other X6 Compute Nodes.

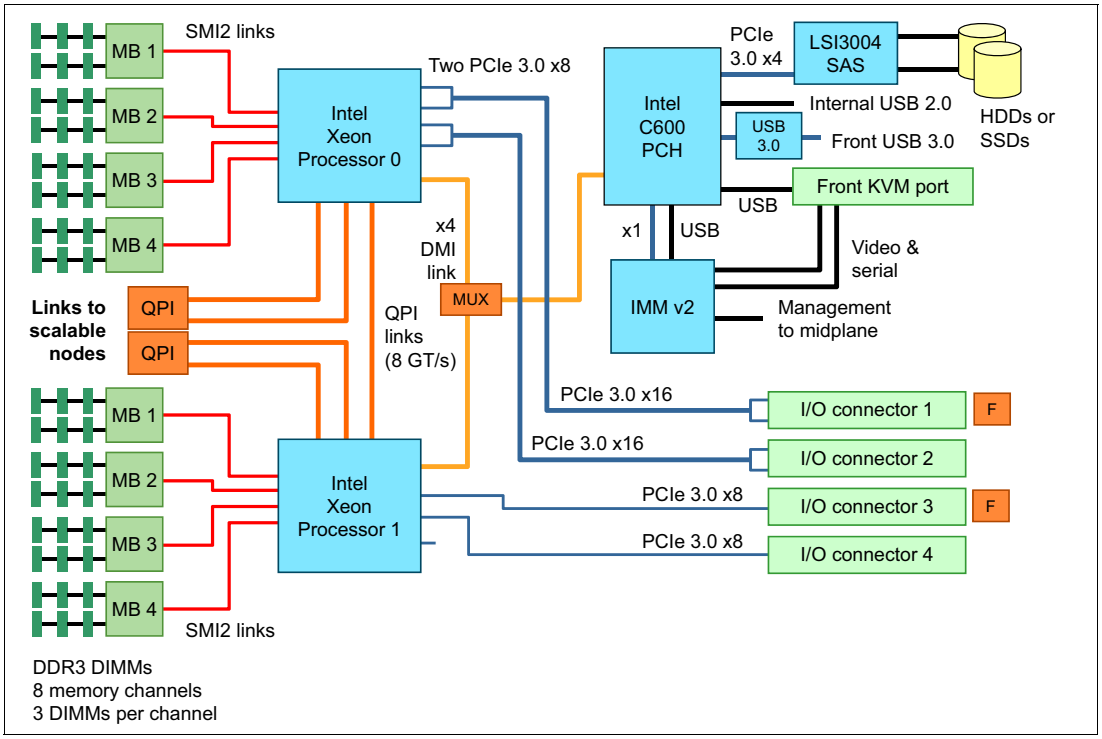

Figure 3-3 2-socket X6 system architecture

The Lenovo Flex System X6 Compute Node has the following standard system architecture features:

- Two 2011-pin type R (LGA-2011) processor sockets.
- An Intel C602j PCIe Controller Hub (PCH).
- Two integrated memory controllers. Each memory controller provides two Intel SMI2 channels that interface with Scalable Memory Buffers.
- Two memory channels per Memory Buffer.
- Up to three DIMMs per memory channel.
- A total of 48 DDR3 DIMM sockets.
- Support for LRDIMMs and RDIMMs.
- One LSI 3004 SAS controller with integrated RAID 0, 1, 10, and 5 to the two 2.5-inch internal drive bays.
- Integrated Management Module II (IMMv2) for systems management.
- Four PCIe 3.0 I/O adapter connectors: Slots 1 and 2 are PCIe 3.0 x16 links composed of two x8 links, and slots 3 and 4 are PCIe 3.0 x8 links.
- Two internal and one external USB connectors.

3.3 Scalability and QPI connectivity

Each X6 Compute Node has two processors. Depending on the processors used, up to four X6 Compute Nodes can be connected to form a single complex:

- Systems with the Intel Xeon E7-4800 v2 can scale to a single 4-socket scalable complex.
- Systems with the Intel Xeon E7-8800 v2 can scale to a single 8-socket scalable complex.
- Systems with the Intel Xeon E7-2800 v2 cannot scale beyond a 2-socket configuration.

This compute node connectivity is done through QPI links at the front of the compute nodes as shown in Figure 3-4.

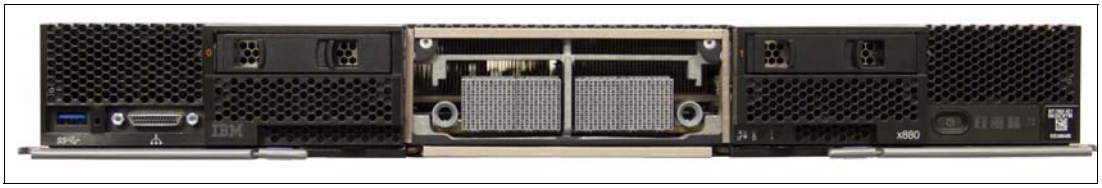

Figure 3-4 X6 Compute Node with the external QPI scalability ports exposed

Processor-to-processor communication is carried over Quick Path Interconnect (QPI) link. Each processor has three QPI links to connect to other processors. The number of QPI links and the connections paths depends on whether the complex is a 2-socket, 4-socket, or 8-socket complex.

2-socket scalability

As a 2-socket (unscaled) server, the two processors are connected through one QPI link. The other two QPI links are not connected. This is shown in Figure 3-5.

Figure 3-5 QPI connections in a 2-socket configuration

4-socket scalability

To form a 4-socket complex, two X6 Compute Nodes are connected (either x880 X6 or x480 X6) using a scalability connector at the front of the server. The Lenovo Flex System x880 X6 4-Socket Scalability Connector shown at Figure 3-6.

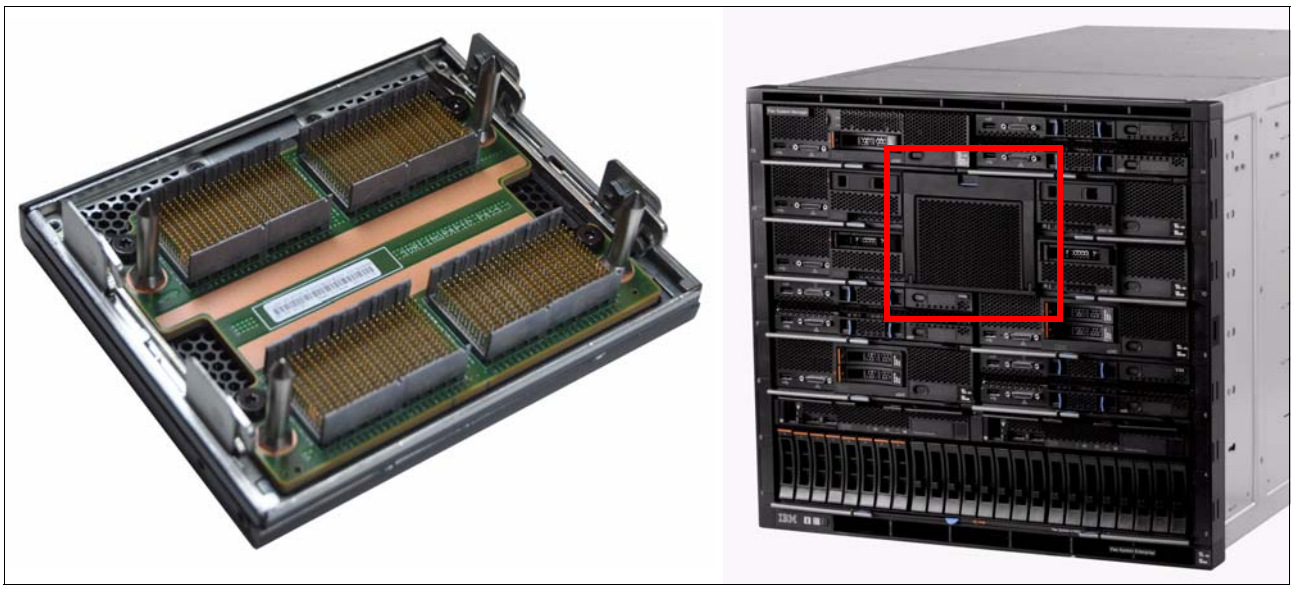

Figure 3-6 Lenovo Flex System x880 X6 4-Socket Scalability Connector

As a 4-socket complex, all three QPI links of each processor are connected in a cross-bar format. Figure 3-7 shows the QPI connectivity; on the right the connections are overlaid on a photo of the inside of the 4-socket scalability connector for reference.

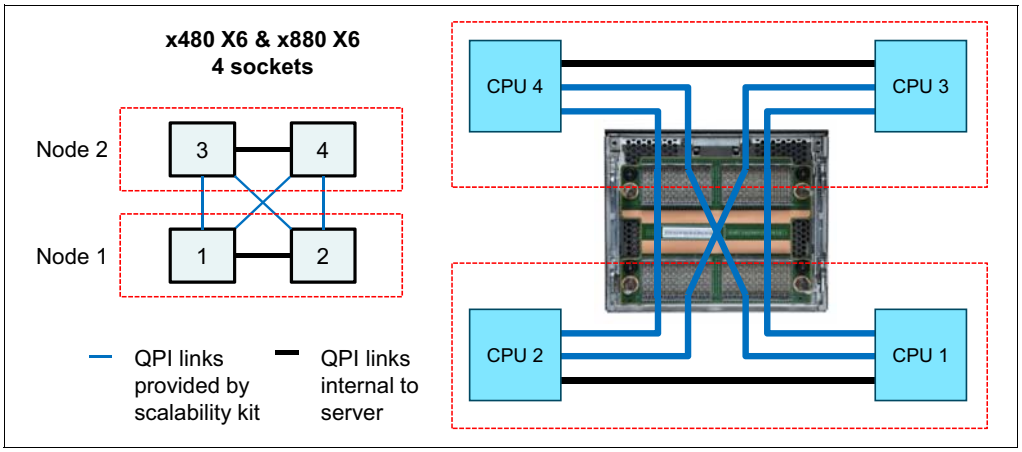

Figure 3-7 QPI connections in a 4-socket configuration

8-socket scalability

The Lenovo Flex System x880 X6 8-Socket Scalability Connector is used to form an 8-socket complex using four X6 Compute Nodes, as shown at Figure 3-8 on page 29.

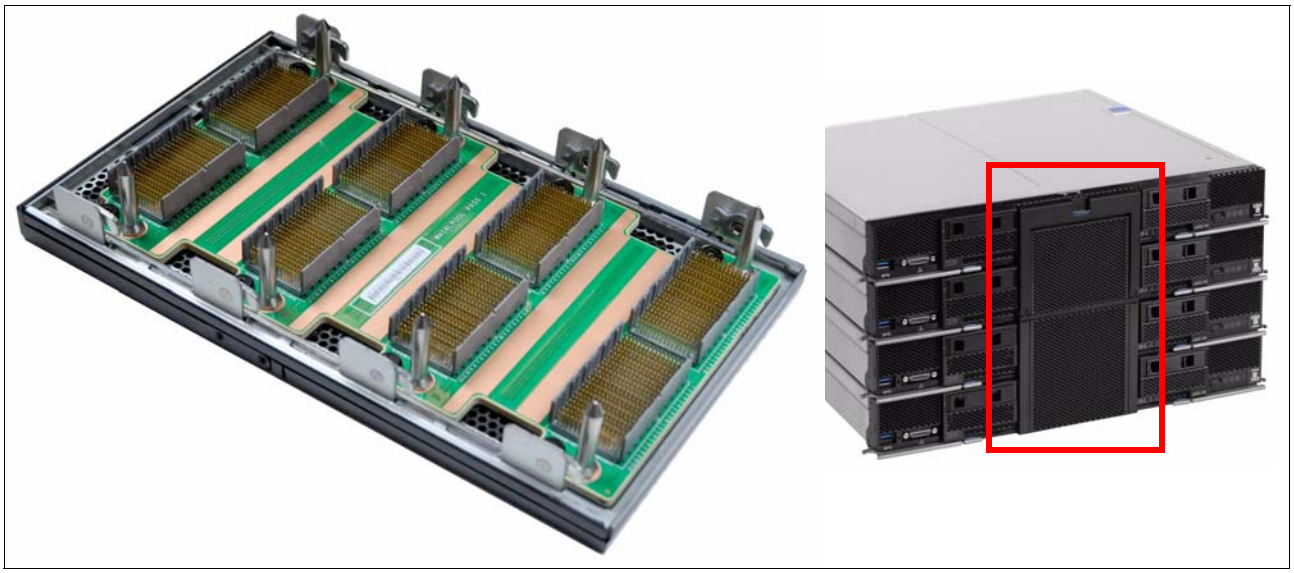

Figure 3-8 Lenovo Flex System x880 X6 8-Socket Scalability Connector

To form an 8-socket complex, four X6 Compute Nodes are connected using an 8-socket scalability connector which connects all three QPI links in each processor. Figure 3-9.

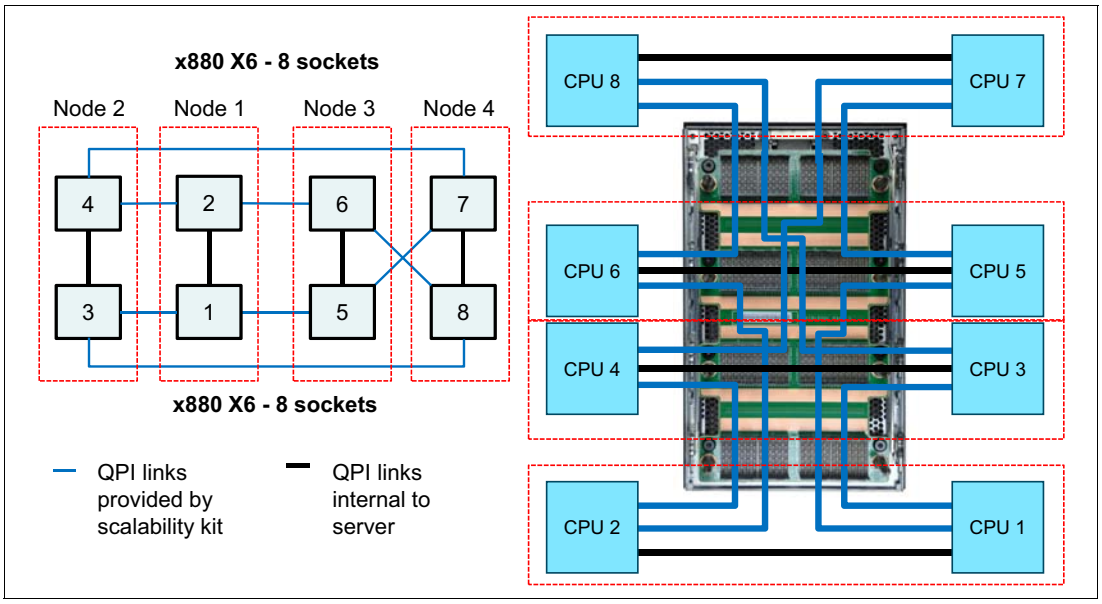

Figure 3-9 QPI connections in an 8-socket configuration

3.4 X6 node resiliency

The X6 Compute Nodes offers fault tolerance at the processor, node, and partition levels:

- Processor failover:

  The 2-socket X6 Compute Node is designed so that it can tolerate the failure of CPU 1 provided there is an operating system boot path from CPU 2. This feature decreases downtime in the event of a processor failure. Downtime is decreased by allowing the machine to boot on another processor until time can be taken to engage Lenovo service.

When the server detects a failure of CPU 1, the IMM reroutes the Intel I/O Hub to CPU 2 instead. The system will boot using CPU 2, albeit with reduced functionality. This failover to a surviving processor is automatic and is handled by the system at boot time.

- Node failover:

  The x480 and x880 can also tolerate a node-failure within a scalable complex. The X6 multi-node scaled complex supports one failing node at a time.

  A multi-node configuration maintains proper partitions and degrades only the partitions that have a failed node. The ability for the system to boot to the OS is dependent upon the OS support and connectivity. Storage media support needs connectivity to the remaining partition resources that are supplied by the degraded partition.

  For more information on scalability and partitioning capabilities, see 5.5, “Scalability and partitioning” on page 100.

- Partition failover:

  A node failover is the last recovery action taken to try to boot a compute node if a catastrophic scenario occurs and results in link failures between the nodes. The system disables QPI, CPU, and memory. If the system has exhausted all of these mechanisms to try to boot a multi-node partition, but the link failure is so severe between the nodes that it cannot be determined if it is present, it performs node failover, so the nodes will reconfigure themselves to act as single systems and boot independently.

  Power failure: A power fault on one of the nodes in a multi-node complex configured as a single whole-system partition will bring down the entire complex. If the complex is configured as independent partitions, a power fault on the primary node will bring down the entire complex, but a power fault on a secondary node will shut down only the partition that it is in, and the primary node remains powered on.

  When a failure occurs in a x480 dual-node complex, the remaining node has the legacy Platform Controller Hub (PCH).

  For an x880 quad-node complex with a single failure, these variations are available:

  - One 8-socket sever with a single failure results in the following 6-socket configurations:
    - 2-3-4 (node 2 has the legacy PCH)
    - 1-3-4 (node 1 has the legacy PCH)
    - 1-2-4 (node 1 has the legacy PCH)
  - Two 4-socket servers with a single failure results in one 4-socket and one 2-socket configurations:
    - 3-4 and 2 (nodes 2 and 3 have the legacy PCH)
    - 3-4 and 1 (nodes 1 and 3 have the legacy PCH)
    - 1-2 and 3 (nodes 1 and 3 have the legacy PCH)
    - 1-2 and 4 (nodes 1 and 4 have the legacy PCH)
  - One 4-socket server and two 2-socket servers with a single failure results in either one 4-socket and one 2-socket or three 2-socket servers:
    - 1-2 and 4 (nodes 1 and 4 have the legacy PCH)
    - 1-2 and 3 (nodes 1 and 3 have the legacy PCH)
    - 1, 3, and 4 (nodes 1, 3, and 4 have the legacy PCH)
    - 2, 3, and 4 (nodes 2, 3, and 4 have the legacy PCH)
  - Four 2-socket servers with a single failure results in three 2-Socket configurations:

3.5 Processor architecture

The X6 Flex Compute nodes use the Intel Xeon processor E7-2800/4800/8800 v2 product family. The Intel Xeon processors that are used in the X6 systems are follow-ons to the Intel Xeon processor E7-2800/4800/8800 product family. New processors feature new Intel microarchitecture (formerly code-named “IvyBridge-EX”) and new 22 nm manufacturing process that provide higher core count, larger cache sizes, higher core frequencies, and higher memory speeds. In addition, these new processors support more memory with up to 24 DIMMs per processor and faster low-latency I/O with integrated PCIe 3.0 controllers.

3.5.1 Processor features

There are three groups of the Intel Xeon processor E7 family that are used in the X6 Flex Compute nodes and support scaling to separate levels:

- The Intel Xeon processor E7-2800 v2 product family is used in the x280 X6. This family supports two-socket configurations.
- The Intel Xeon processor E7-4800 v2 product family is used in the x480 X6 to scale to four-socket configurations.
- The Intel Xeon processor E7-8800 v2 product family is used in the x880 X6 to scale to eight-socket configurations.

The X6 systems support the latest generation of Intel Xeon processor E7-2800 v2, E7-4800 v2 and E7-8800 v2 product families which offer the following key features:

- Up to 15 cores and 30 threads (using Hyper-Threading feature) per processor
- Up to 37.5 MB of L3 cache
- Up to 3.4 GHz core frequencies
- Up to 8 GT/s bandwidth of QPI links
- Integrated memory controller with four SMI2 channels that support up to 24 DDR3 DIMMs
- Up to 1600 MHz DDR3 memory speeds and up to 2667 MHz SMI link speeds
- Integrated PCIe 3.0 controller with 32 lanes per processor
- Intel Virtualization Technology (VT-x and VT-d)
- Intel Turbo Boost Technology 2.0
- Intel Advanced Vector Extensions (AVX)
- Intel AES-NI instructions for accelerating of encryption
- Advanced QPI and memory reliability, availability, and serviceability (RAS) features
- Machine Check Architecture recovery (non-execution and execution paths)
- Enhanced Machine Check Architecture Gen1
- Machine Check Architecture I/O
- Security technologies: OS Guard, Secure Key, Intel TXT

The new Intel Xeon processors feature advanced technologies described below.

Intel Advanced Encryption Standard - New Instructions

Advanced Encryption Standard (AES) is an encryption standard that is widely used to protect network traffic and sensitive data. Advanced Encryption Standard - New Instructions (AES-NI), available with the E7 processors, implements certain complex and performance intensive steps of the AES algorithm by using processor hardware. AES-NI can accelerate the performance and improve the security of an implementation of AES over an implementation that is completely performed by software.

For more information, see the Intel Advanced Encryption Standard (AES) New Instructions (AES-NI) web page:

http://intel.ly/YeE7QS

Intel Virtualization Technology

Intel Virtualization Technology (Intel VT) is a suite of processor and I/O hardware enhancements that assists virtualization software to deliver more efficient virtualization solutions and greater capabilities.

Intel Virtualization Technology for x86 (Intel VT-x) allows the software hypervisors to better manage memory and processing resources for virtual machines (VMs) and their guest operating systems.

Intel Virtualization Technology for Directed I/O (Intel VT-d) helps improve I/O performance and security for VMs by enabling hardware-assisted direct assignment and isolation of I/O devices.

For more information, see the web page about Intel Virtualization Technology:

http://www.intel.com/technology/virtualization

Hyper-Threading Technology

Intel Hyper-Threading Technology enables a single physical processor to run two separate code streams (threads) concurrently. To the operating system, a processor core with Hyper-Threading is seen as two logical processors. Each processor has its own architectural state, that is, its own data, segment, and control registers, and its own advanced programmable interrupt controller (APIC).

Each logical processor can be individually halted, interrupted, or directed to run a specified thread, independently from the other logical processor on the chip. The logical processors share the execution resources of the processor core, which include the execution engine, the caches, the system interface, and the firmware.

Hyper-Threading Technology is designed to improve server performance. This process is done by using the multi-threading capability of operating systems and server applications in such a way as to increase the use of the on-chip execution resources available on these processors. Application types that make the best use of Hyper-Threading are virtualization, databases, email,, and web servers.

For more information, see the Intel Hyper-Threading Technology web page:

http://www.intel.com/technology/platform-technology/hyper-threading

Turbo Boost Technology 2.0

The Intel Xeon E7-8800/4800/2800 v2 processor family brings enhanced capabilities of changing processor speed with new generation of Intel Turbo Boost 2.0 technology.

Intel Turbo Boost Technology dynamically saves power on unused processor cores and increases the clock speed of the cores in use. Depending on current workload Intel Turbo Boost Technology allows a dynamic increase in the clock speed of the active cores to gain a performance boost. For example, a 3.4 GHz 15-core processor can overclock the cores up to 3.7 GHz.

Turbo Boost Technology is available on a per-processor basis for the X6 systems. For ACPI-aware operating systems and hypervisors such as Microsoft 2008/2012, RHEL 5/6, SLES 11 and VMware ESXi 4.1 and later, no changes are required to take advantage of it.

Turbo Boost Technology can be engaged with any number of cores enabled and active, resulting in increased performance of both multi-threaded and single-threaded workloads.

Turbo Boost Technology dynamically saves power on unused processor cores and increases the clock speed of the cores in use. In addition, it can temporarily increase speed of all cores by intelligently managing power and thermal headroom. For example, a 2.5 GHz 15-core processor (E7-8880 v2) can temporarily run all 15 active cores at 2.9 GHz. With only two cores active, the same processor can run those active cores at 3.0 GHz. When the cores are needed again, they are turned back on dynamically and the processor frequency is adjusted accordingly.

When temperature, power, or current exceeds factory-configured limits and the processor is running above the base operating frequency, the processor automatically steps the core frequency back down to reduce temperature, power, and current. The processor then monitors temperature, power, and current and re-evaluates. At any given time, all active cores run at the same frequency.

For more information, see the Intel Turbo Boost Technology web page:

http://www.intel.com/technology/turboboost/

QuickPath Interconnect

The Intel Xeon E7 processors implemented in Lenovo X6 servers include two integrated memory controllers in each processor. Processor-to-processor communication is carried over shared-clock or coherent Intel QuickPath Interconnect (QPI) links. Each processor has three QPI links to connect to other processors.

We describe how the QPI links are used to scale the X6 Compute Nodes in 3.3, “Scalability and QPI connectivity” on page 27.

For more information about QPI, see the Intel QuickPath web page:

http://www.intel.com/technology/quickpath

Intel Data Direct I/O

For I/O, Intel no longer has a separate I/O hub. It now integrates PCI Express 3.0 I/O into the processor itself. Data Direct I/O helps to optimize data transfer between local CPU and PCIe devices. The combination of Data Direct I/O and PCI 3.0 provides a significantly higher I/O performance with lower latencies and reduced power consumption.

For more information, see the Intel Data Direct I/O Technology web page:

http://www.intel.com/content/www/us/en/io/direct-data-i-o.html

RAS features

The Intel Xeon processor E7 family of processors has additional reliability, availability, and serviceability (RAS) features on their interconnect links (SMI and QPI):

- Cyclic redundancy checking (CRC) on the QPI links

  The data on the QPI link is checked for errors.

- QPI packet retry

  If a data packet on the QPI link has errors or cannot be read, the receiving processor can request that the sending processor try resending the packet.

- QPI clock failover

  If there is a clock failure on a coherent QPI link, the processor on the other end of the link can become the clock. This action is not required on the QPI links from processors to I/O hubs, as these links are asynchronous.

- QPI self-healing

  If there are persistent errors detected on a QPI link, the link width can be reduced dynamically to allow the system to run in a degraded mode until repair can be performed. QPI link can reduce its width to a half width or a quarter width, and slowdown its speed.

- Scalable memory interconnect (SMI) packet retry

  If a memory packet has errors or cannot be read, the processor can request that the packet be resent from the memory buffer.

Machine Check Architecture recovery

The Intel Xeon processor E7 family also features Machine Check Architecture (MCA) recovery, a RAS feature that enables the handling of system errors that otherwise require that the operating system be halted. For example, if a dead or corrupted memory location is discovered, but it cannot be recovered at the memory subsystem level, and provided it is not in use by the system or an application, an error can be logged and the operation of the server can continue. If it is in use by a process, the application to which the process belongs can be stopped or informed about the situation.

Implementation of the MCA recovery requires hardware support, firmware support (such as found in the UEFI), and operating system support. Microsoft, SUSE, Red Hat, VMware, and other operating system vendors include or plan to include support for the Intel MCA recovery feature on the Intel Xeon processors in their latest operating system versions.

New MCA recovery features of the Intel Xeon processor E7-2800/4800/8800 v2 product family include:

- Execution path recovery: Ability to work with hardware and software to recognize and isolate the errors that were delivered to the execution engine (core).
- Enhanced MCA (eMCA) Generation 1: Provides enhanced error log information to the operating system, hypervisor, or application that can be used to provide better diagnostic and predictive failure analysis for the system. This enables higher levels of uptime and reduced service costs.

Security improvements

The Intel Xeon E7-2800/4800/8800 v2 processor family has several important security improvements which help to protect systems from different types of security threats.

- Intel OS Guard: Evolution of Intel Execute Disable Bit technology which helps to protect from escalation of privilege attacks by preventing code execution from user space memory pages while in kernel-mode. It helps to protect from certain types of malware attacks.
- Intel Trusted Execution Technology (Intel TXT), Intel VT-x and Intel VT-d: New hardware-based techniques which allows to isolate virtual machines and boot VMs only in trusted environment. In additional malware infected VMs cannot affect another VMs on the same host.
- Intel Secure Key: Provides hardware random numbers generation without storing any data in system memory. It keeps generated random numbers out of sight of malware and therefore enhances encryption protection.

For more information, read the following Intel white paper titled *Crimeware Protection: 3rd Generation Intel Core vPro Processors*:

http://intel.ly/1py6l3u

3.6 Memory architecture

The Lenovo Flex System X6 Compute Nodes support DDR3 memory with ECC protection. The X6 supports up to 48 DIMMs when two processors are installed (24 DIMMs per processor), and up to 192 DIMMs with 8-socket scaled configuration.

Topics in this section:

Each processor has two integrated memory controllers, and each memory controller has two Scalable Memory Interconnect generation 2 (SMI2) links that are connected to two scalable memory buffers. Each memory buffer has two DDR3 channels, and each channel supports three DIMMs, for a total of 24 DIMMs per processors.

Figure 3-10 show the processor’s memory architecture.

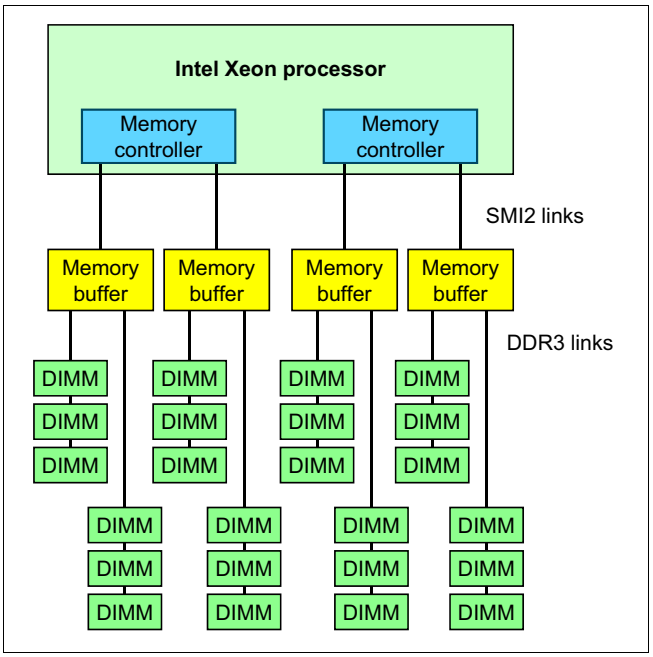

Figure 3-10 Intel Xeon processor E7 v2 memory architecture

3.6.1 Operational modes

There are two memory modes that are supported by the Intel Xeon processor E7-2800 v2/4800 v2/8800 v2 product families

- Performance mode

  In this operation mode each DDR3 channel works independently and it is addressed individually by burst lengths of 8 bytes (64 bits). The Intel SMI2 channel operates at twice the DDR3 speed. All channels can be populated in any order and modules have no matching requirements. All memory modules may run at different DIMM timings, but all DIMMs must run at the same speed.

  Chipkill (also known as Single Device Data Correction or SDDC) is supported in Performance mode. Redundant Bit Steering (RBS) is not supported.

  Although DIMMs can be populated in any order in this mode, for best performance, memory modules should be placed based on round robin algorithm between SMI2 channels and alternating between DDR channels. For more information about DIMMs population order, see 4.8, “Memory population order” on page 57.

- RAS (Lockstep) mode

  RAS stands for reliability, availability, and serviceability. In this operation mode (also known as Lockstep mode), the memory controller operates two DDR3 channels behind one memory buffer as single channel.

  In RAS mode the SMI2 channel operates at the DDR3 transfer rate. DIMMs must be installed in pairs.

  Since data is moved using both channels at once, additional advanced memory protection schemes can be implemented to provide protection against both single-bit and multi-bit errors:

  – Chipkill memory, also known as Single Device Data Correction (SDDC)
  – Redundant Bit Steering (RBS)

  The combination of these two RAS features is also known as Double Device Data Correction (DDDC).

Mirroring and sparing are also supported in both modes as described in 3.6.2, “Memory mirroring” on page 38 and 3.6.3, “Rank sparing” on page 39 respectively.

Figure 3-11 shows the two modes. In RAS mode, both channels off one memory buffer are in lockstep with each other.

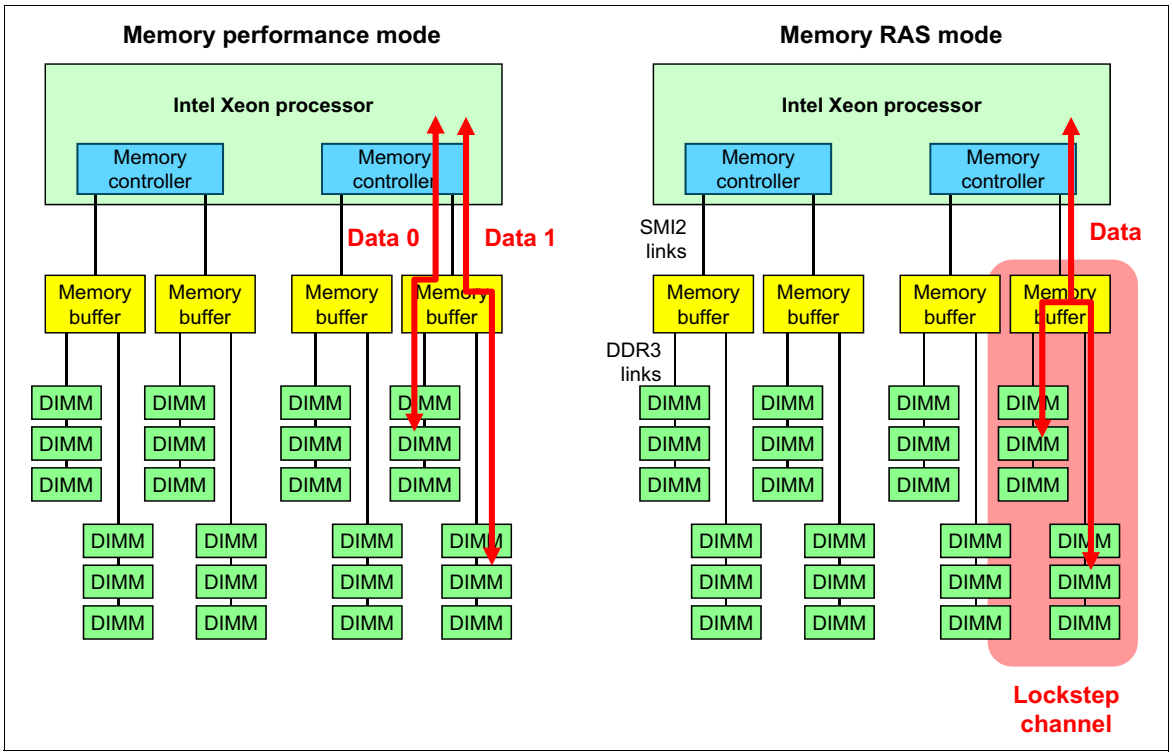

Figure 3-11 Memory modes - Performance mode (left) and RAS mode (right)

Table 3-1 shows speed and maximum bandwidth for the SMI2 and DDR3 channels in both modes.

Table 3-1 Memory mode comparison

Bus	Frequency	Transfers/ clock	Transfer rate	Data width	Channel bandwidth
Performance mode					
SMI2 Bus	1333 MHz	2	2666 MT/s	64 bit	21.3 GBps
DDR3 Bus	667 MHz	2	1333 MT/s	64 bit	10.7 GBps
RAS mode					
SMI2 Bus	800 MHz	2	1600 MT/s	64 bit	12.8 GBps
DDR3 Bus	800 MHz	2	1600 MT/s	64 bit	12.8 GBps

In addition to Performance and RAS modes, the memory subsystem has additional reliability, availability, and serviceability (RAS) features that can be enabled from the Unified Extensible Firmware Interface (UEFI):

- Memory mirroring
- Rank sparing

3.6.2 Memory mirroring

To improve memory reliability and availability the memory controller can mirror memory data to two memory channels. To enable mirroring feature, you must have both memory channels of a processor populated with the same DIMM type and amount of memory.

Memory mirroring provides the user with a redundant copy of all code and data addressable in the configured memory map. Two copies of the data are kept, similar to the way RAID-1 writes to disk. Reads are interleaved between memory channels. The system automatically uses the most reliable memory channel as determined by error logging and monitoring.

If errors occur, only the alternative memory channel is used until bad memory is replaced. Because a redundant copy is kept, mirroring results in only half the installed memory being available to the operating system. Memory mirroring does not support asymmetrical memory configurations and requires that each channel is populated in identical fashion. For example, you must install two 4 GB of identical memory modules equally and symmetrically across the two memory channels to achieve 4 GB of mirrored memory.

Memory mirroring is independent of the operating system. There is a slight memory performance trade-off when memory mirroring is enabled.

Memory mirroring feature can be used in conjunction with performance or RAS modes:

- When performance mode is used, memory mirroring duplicates data between memory channels of the two memory buffers connected to one memory controller.
- In RAS (Lockstep) mode memory mirroring duplicates data between memory buffers connected to the same memory controller.

Figure 3-12 shows how memory mirroring is implemented in Performance mode (left) and RAS mode (right).

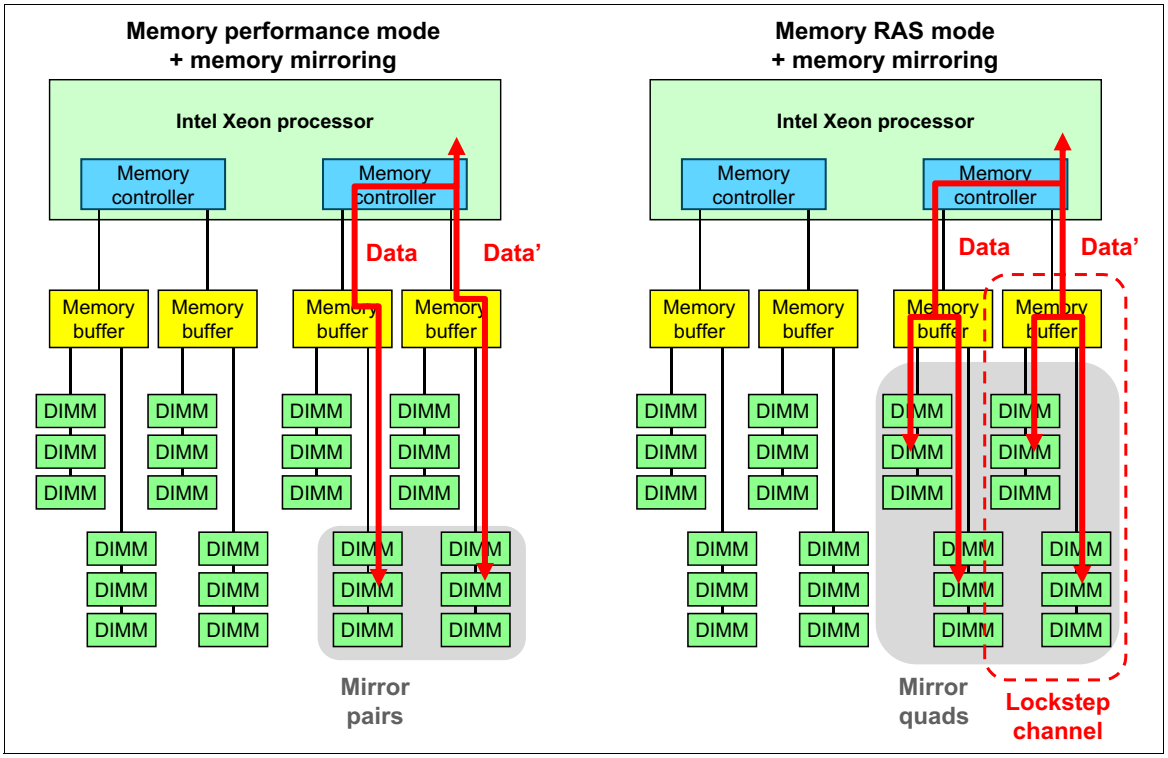

Figure 3-12 Memory mirroring with used with Performance mode (left) and RAS mode (right)

Memory mirroring rules:

- The server supports single-socket memory mirroring. The compute book memory channel 0 mirrors memory channel 1, and memory channel 2 mirrors memory channel 3. This mirroring provides redundancy in memory but reduces the total memory capacity in half.
- DIMMs must be installed in pairs for each compute book when using the memory mirroring feature.
- The DIMM population must be identical (size, organization, etc.) for memory channel 0 and memory channel 1, and identical for memory channel 2 and memory channel 3.
- Memory mirroring reduces the maximum available memory by half of the installed memory. For example, if the server has 64 GB of installed memory, only 32 GB of addressable memory is available when memory mirroring is enabled.

3.6.3 Rank sparing

In rank-sparing mode, one rank is held in reserve as a spare of the other ranks in the same memory channel. There are eight memory channels per processor.

Memory rank sparing provides a degree of redundancy in the memory subsystem, but not to the extent of mirroring. In contrast to mirroring, sparing leaves more memory for the operating system. In sparing mode, the trigger for failover is a preset threshold of correctable errors. When this threshold is reached, the content is copied to its spare rank. The failed rank is then taken offline, and the spare counterpart is activated for use.

In rank sparing mode, one rank per memory channel is configured as a spare. The spare rank must have identical or larger memory capacity than all the other ranks (sparing source ranks) on the same channel.

For example, If dual-rank DIMMs are installed, all of the same capacity, then there are 6 ranks total for each memory channel (three DIMMs per channel). This means that 1 of the 6 ranks are reserved and 5 of the 6 are usable for operating system use.

Memory sparing is independent of the operating system. There is a slight memory performance trade-off when memory sparing is enabled.

Rank sparing feature can be used in addition to performance or RAS modes:

- When performance mode is used, rank sparing duplicates data between memory modules of same channel of one memory buffer. In the event of an imminent failure (red X in Figure 3-13 on page 40), that rank is taken offline and the data is copied to the spare rank.
- When RAS (Lockstep) mode is used, rank sparing duplicates data between memory channels of one memory buffer. In the event of an imminent failure (red X in Figure 3-13 on page 40), that rank is taken offline and the data is copied to the spare rank. In addition, the partner rank on the other channel that is connected to the same memory buffer is also copied over.

Figure 3-13 shows the rank sparing usage in conjunction with Performance mode (left) and RAS mode (right).

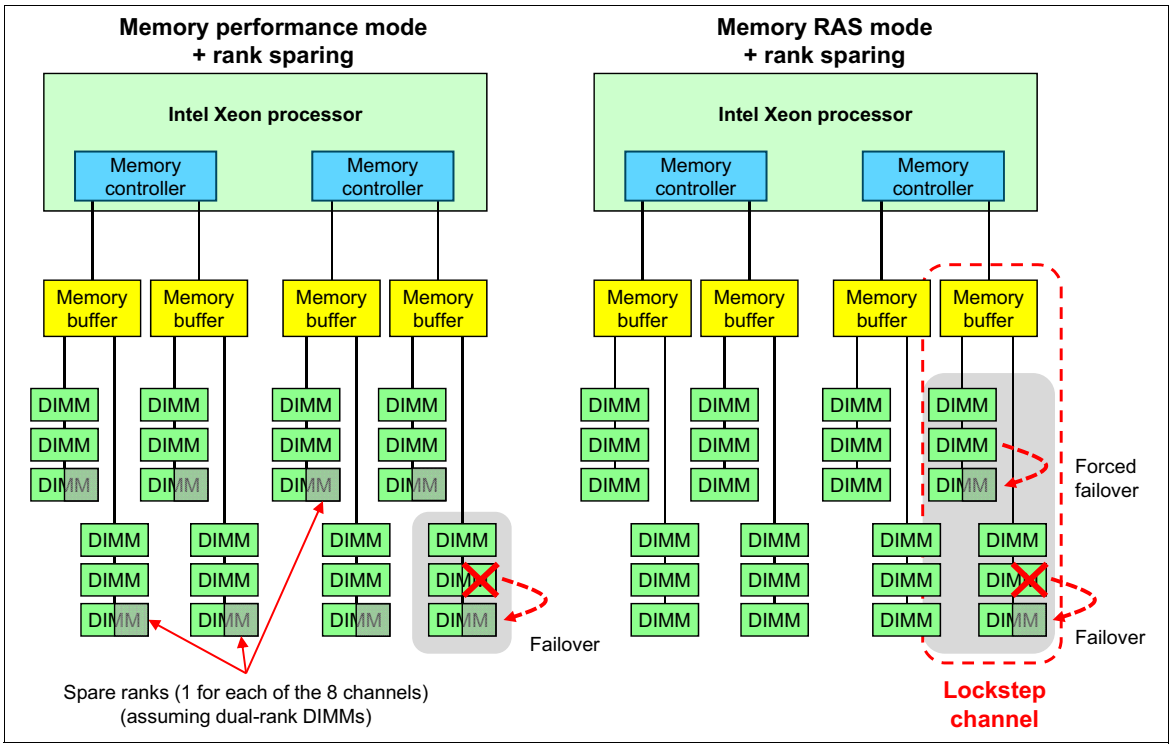

Figure 3-13 Rank sparing: Performance mode (left) and RAS mode (right)

Configuration rules for rank sparing:

- Memory rank sparing is not supported if memory mirroring is enabled.
- The spare rank must have identical or larger memory capacity than all the other ranks on the same DDR3 channel.
- When single-rank DIMMs are used, a minimum of two rank DIMMs must be installed per memory channel to support memory sparing.
- When multi-rank DIMMs are used, one multi-rank DIMM can be installed per memory channel to support memory sparing.
- The total memory available in the system is reduced by the amount of memory allocated for the spare ranks.

3.6.4 Chipkill memory technology

Chipkill memory technology, an advanced form of error correction code (ECC), is available for the X6 servers. Chipkill (also known as Single Device Data Correction, or SDDC) protects the memory in the system from any single memory chip failure. It also protects against multibit errors from any portion of a single memory chip.

Chipkill on its own is able to provide 99.94% memory availability to the applications without sacrificing performance with standard ECC DIMMs.

Chipkill is used in both Performance mode and RAS mode.

3.6.5 Redundant bit steering

Redundant bit steering provides the equivalent of a hot-spare drive in a RAID array. It is based in the memory controller and senses when a chip on a DIMM fails and when to route the data around the failed chip.

Redundant bit steering is only used in RAS mode.

Within the system, the models of the X6 servers using the Intel Xeon processor E7 v2 family support the Intel implementation of Chipkill plus redundant bit steering, which Intel calls Double Device Data Correction (DDDC).

Redundant bit steering uses the ECC coding scheme that provides Chipkill coverage for x4 DRAMs. This coding scheme leaves the equivalent of one x4 DRAM spare in every pair of DIMMs. If a chip failure on the DIMM is detected, the memory controller can copy data from the failed chip through the spare x4.

Redundant bit steering operates automatically without issuing a Predictive Failure Analysis (PFA) or light path diagnostics alert to the administrator, although an event is logged to the service processor log. In RAS (Lockstep) mode, after the second DRAM chip failure on the DIMM, additional single bit errors will result in PFA and light path diagnostics alerts.

3.6.6 Advanced Page Retire

Advanced Page Retire is an algorithm to handle memory errors. It is a built-in sophisticated error handling firmware that leverages and coordinates memory recovery features, balancing the goals of maximum up time and minimum repair actions.

The algorithm uses short- and long-term thresholds per memory rank with leaky bucket and automatic sorting of memory pages with the highest correctable error counts. First, it leverages hardware recovery features, followed by software recovery features, to optimize recovery results for both newer and older operating systems and hypervisors.

When recovery features are exhausted, firmware will issue a Predictive Failure Alert. Memory which has failed completely will be held offline during reboots until repaired. Failed DIMMs are indicated by light path diagnostics LEDs physically at the socket location.

Lenovo performs thorough testing to verify the features and coordination between the firmware and the operating system or hypervisor.

3.7 Internal Storage

In this section we give an overview of different technologies used in the storage subsystem of the X6 node family.

The X6 Compute Node has two 2.5-inch hot-swap drive bays that are accessible from the front of the server, as shown in Figure 4-2 on page 47. These bays are connected to the integrated 4-port LSI SAS 3004 12 Gbps SAS/SATA RAID-on-Chip (ROC) controller.

The integrated LSI SAS 3004 ROC includes the following features:

- Four-port controller with 12 Gbps throughput per port
- PCIe x4 2.0 host interface
- Two SAS ports that are routed internally to the two hot-swap drive bays
- Supports RAID levels 0, 1, 10 and 1E; optional RAID 5 and RAID 50.

RAID 5 and RAID 50 are supported by a Features on Demand license upgrade. See Chapter 4, “Flex System X6 Compute Node” on page 45 for details.

3.8 PCIe 3.0

The X6 family of compute nodes supports the new generation PCI Express (PCIe) protocol, Version 3.0. PCIe 3.0 is the evolution of the PCI Express I/O standard, which brings double the bandwidth of PCIe 2.0 while preserving compatibility with previous generations of the PCIe protocol. This means that PCIe 1.x and 2.x generation cards will work in PCIe 3.0-capable slots, and PCIe 3.0 cards will work when plugged into PCIe slots from previous generations.

PCI 3.0 uses a 128b/130b encoding scheme that is more efficient than 8b/10b encoding, which is used in PCIe 2.0. Therefore, it doubles the bandwidth at 8 GT/s speeds.

Each Intel Xeon E7-2800/4800/8800 v2 processor contains an Integrated I/O (IIO) module that provides 32 lanes of PCIe 3.0. These 32 lanes can be split into any combination of x4, x8, and x16.

Table 3-2 shows a comparison of PCIe capabilities of the previous eX5 architecture to the X6 architecture.

Table 3-2 Lenovo eX5 and X6 PCIe capabilities

Specification	X5™ generation	X6 generation
PCIe version	PCIe 2.0	PCIe 3.0
Raw bit rate	5 GT/s	8 GT/s
Encoding	8b/10b	128b/130b
Bandwidth per lane	~500 MB/s	~ 1000 MB/s

For more information on PCIe 3.0, see the PCI Express Base 3.0 specification by PCI-SIG:

http://www.pcisig.com/specifications/pciexpress/base3/

3.9 Partitioning

The Lenovo Flex System X6 Compute Nodes support Flex Node Partitioning, which allows for a multi-node system to be configured as multiple smaller node systems without physically reconfiguring.

The following partitioning configurations are supported:

- Single-node configuration (x280 X6, x480 X6, or x880 X6):
  - One 2-socket server
- Two-node configuration (x480 X6 or x880 X6):
  - One 4-socket server
  - Two 2-socket servers

- Four-node configuration (x880 X6 only):
  - One 8-socket server
  - Two 4-socket servers
  - One 4-socket server and two 2-socket servers
  - Four 2-socket servers

There are a few things to consider for partitioning:

- The foremost requirement is for all nodes to boot so that they are statically configured at the same time. The complex is not dynamically capable of switching configurations without all of the nodes coming down to reconfigure the QPI links during the boot of the system firmware.
- There are configurations for the multi-node configurations where the nodes of a partition need to be in adjacent slots. The end nodes are considered adjacent.
- After the initial insertion, if the node is removed, slots changed, flash updated, or chassis changed, the system should maintain its previous partition status.
- When a multi-node X6 configuration is initially installed after a change in the number of nodes in a complex or a replacement of one of the nodes, the default behavior should be to create a system that encompasses all nodes.
- The node with the lowest slot ID is the primary node and the rest are the subordinate nodes. The subordinate nodes will receive the configuration information from the primary node through the IMMv2.
- The primary node is responsible for disseminating the information to the subordinates for partitioning boot.

3.10 UEFI

The Unified Extensible Firmware Interface (UEFI) is a pre-boot environment that provides an interface between server firmware and the operating system. It replaces BIOS as the software that manages the interface between server firmware, operating system, and hardware initialization and eliminates the 16-bit, real-mode limitation that BIOS had.

These are the main advantages of UEFI:

- 64-bit code architecture, no restrictions for code size
- Modular design, so adapter vendors can add any functions in ROM
- Adapter configuration can move into F1 Setup
- No beep codes, all errors covered by light path diagnostics
- Human-readable Event logs in F1 Setup
- UEFI settings can be accessed out of band by using the Lenovo Advanced Setup Utility (ASU) through the IMM web interface
- UEFI code updates can be performed in band through the operating system and out of band through the IMM web interface

Get more information on the UEFI website:

http://www.uefi.org/home/

3.11 Integrated Management Module

The Integrated Management Module II (IMM2) is the next generation of the IMM service processor that provides advanced control and monitoring features to manage the X6 Compute Node. IMM2 enables easy console redirection with text and graphics and keyboard and mouse support over the system management LAN connections.

This feature allows the user to display server activities from power-on to full operation remotely, with remote user interaction at virtually any time.

IMM2 monitors the following components:

- System voltages
- System temperatures
- Fan speed control
- Fan tachometer monitor
- Good Power signal monitor
- System ID and planar version detection
- System power and reset control
- Non-maskable interrupt (NMI) detection (system interrupts)
- SMI detection and generation (system interrupts)
- Serial port text console redirection
- System LED control (power, HDD, activity, alerts, and heartbeat)

IMM2 provides these features:

- An embedded web server, which gives you remote control from any standard web browser with Oracle Java plug-in installed
- A command-line interface (CLI), which the administrator can use from a Telnet or SSH session
- Secure Sockets Layer (SSL) and Lightweight Directory Access Protocol (LDAP)
- Built-in LAN and serial connectivity that support virtually any network infrastructure
- Multiple alerting functions to warn systems administrators of potential problems through email, IPMI platform event traps (PETs), and Simple Network Management Protocol (SNMP)

IMM2 has many enhancements when compared to previous generation of IMM:

- Faster processor and more memory, which deliver a more responsive user interface
- Improved system power-on and boot time
- New web user interface that provides a common design across all Lenovo System x products
- Increased color depth and resolution for remote presence function
- IMM2 configuration changes require no reboot and become effective immediately
- Hardware management of non-volatile storage
- 1 Gb Ethernet management capability
- More detailed information for UEFI-detected events
- Simplified update and flashing mechanism
- Security enhancements, including authentication policies

For more detailed information, see the *Integrated Management Module II User's Guide*:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086346

4

Flex System X6 Compute Node

The Lenovo Flex System X6 Compute Node, machine type 7903, is a high-performance compute node designed to take on the most demanding workloads. The performance, flexibility, and resiliency features enable the X6 to run mission-critical workloads, such as key business applications, database, analytics, and large virtual machine deployments.

This chapter covers the following topics:

4.1 Introduction

The Lenovo Flex System X6 Compute Node family includes the three X6 Compute Nodes: x280 X6, x480 X6, and x880 X6. These three servers are all based on a two-socket double-wide compute node and differ only by the processor type installed and the degree to which they can scale up:

- The Flex System x880 X6 Compute Node uses Intel Xeon E7-8800 v2 processors, can scale up to 8 sockets, and supports 8-socket (composed of four compute nodes connected together), 4-socket (composed of two compute nodes connected), and 2-socket (one compute node) configurations.
- The Flex System x480 X6 Compute Node uses Intel Xeon E7-4800 v2 processors, can scale up to 4-socket, and supports 4-socket and 2-socket configurations.
- The Flex System x280 X6 Compute Node uses Intel Xeon E7-2800 v2 processors and supports only 2-socket configurations.

Figure 4-1 shows two x880 X6 Compute Nodes connected to form a 4-socket server.

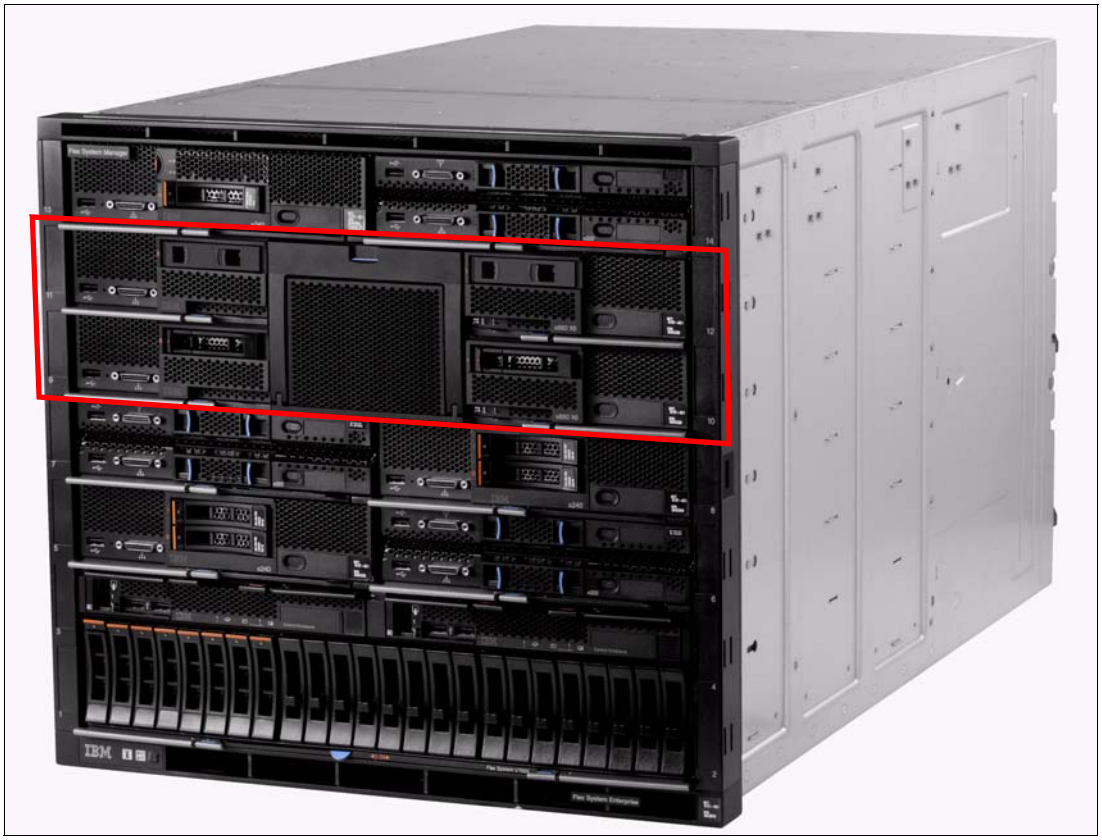

Figure 4-1 Two Flex System x880 Compute Nodes connected to form a 4-socket complex

Figure 4-2 shows the front of the compute node, which includes the location of the controls, LEDs, and connectors. The light path diagnostic panel is on the upper edge of the front panel bezel, in the same place as on the other x86 compute nodes.

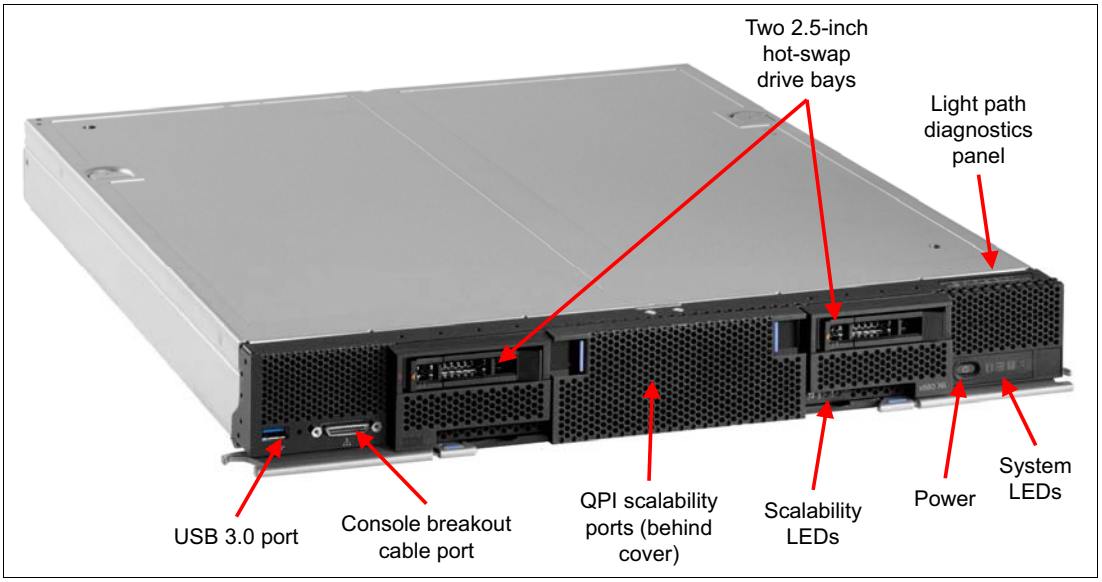

Figure 4-2 The Lenovo Flex System X6 Compute Node

4.2 Specifications

Table 4-1 lists the features of the X6 nodes.

Table 4-1 Lenovo Flex System X6 node specifications

Components	Specification (per node, except where noted)
Form factor	Double-wide compute node.
Chassis support	Lenovo Flex System Enterprise Chassis.
Scalability	x880 X6: ► Scales up to 4-socket by adding one x880 Compute Node + 4S scalability kit ► Scales up to 8-socket by adding three x880 Compute Nodes + 8S scalability kit x480 X6: ► Scales up to 4-socket by adding one x480 Compute Node + 4S scalability kit x280 X6: ► Does not scale

Components	Specification (per node, except where noted)
Processor	Processor selection determines whether the compute node is an x880 X6, x480 X6, or x280 X6. The x880 X6 and x280 X6 are currently CTO only: ► x880 X6 (CTO only): Two Intel Xeon E7-8800 v2 processors, each with 15 cores (up to 2.8 GHz), 12 cores (up to 3.0 GHz), 10 cores (up to 3.2 GHz), or 6 cores (up to 3.4 GHz). Three QPI links. ► x480 X6 (standard models or CTO): Two Intel Xeon E7-4800 v2 processors, each with 15 cores (up to 2.8 GHz), 12 cores (up to 2.6 GHz), 10 cores (2.2 GHz), 8 cores (2.0 GHz), or 6 cores (1.9 GHz). Three QPI links. ► x280 (CTO only): Two Intel Xeon E7-2800 v2 processors, each with 15 cores (up to 2.8 GHz) or 12 cores (2.3 GHz). Two QPI links (1 connected). All processors: QPI link up to 8.0 GTps. Up to 1600 MHz memory speed. Up to 37.5 MB L3 cache per processor. Intel C602J chipset.
Chipset	Intel C602J.
Memory	Up to 48 DIMM sockets (24 DIMMs per processor) using Low Profile (LP) DDR3 DIMMs. RDIMMs and LRDIMMs are supported. 1.5 V and low-voltage 1.35 V DIMMs are supported. Support for up to 1600 MHz memory speed, depending on the processor. Four SMI2 channels to memory buffer per processor. Two memory channels per memory buffer. Supports three DIMMs per channel.
Memory maximums	With LRDIMMs: Up to 3 TB with 48x 64 GB LRDIMMs and two processors. With RDIMMs: Up to 768 GB with 48x 16 GB RDIMMs and two processors.
Memory protection	ECC, Chipkill (for x4-based memory DIMMs), memory mirroring, and memory rank sparing.
Disk drive bays	Standard: Two 2.5-inch hot-swap SAS/SATA drive bays that support SAS, SATA, and SSD drives.
Maximum internal storage	With two 2.5-inch hot-swap drives: Up to 3.2 TB with 1.6 TB 2.5" NL SAS HDDs, or up to 2.4 TB with 1.2 TB 2.5" SAS HDDs, or up to 2 TB with 1 TB 2.5" NL SATA HDDs, or up to 3.2 TB with 1.6 TB 2.5" SAS SSDs. Intermix of SAS and SATA HDDs and SSDs is supported.
RAID support	RAID 0, 1, 1E, 10 with integrated ServeRAID™ M1200e (LSI SAS3004), upgradeable to RAID 5, 50. Two adjacent compute nodes, each with two drives, can form a 4-drive RAID-10 array.
Network interfaces	None standard. Optional Ethernet adapters.
PCI Expansion slots	Up to four I/O connectors for adapters. Two connectors have a PCI Express 3.0 x24 (x16 + x8) interface with support for dual-ASIC adapters (slots 1 and 2). Two other connectors have PCI Express 3.0 x8 interface (slots 3 and 4).
Ports	External: One USB 3.0. Console breakout cable port that provides local KVM and serial ports (cable standard with Flex System chassis, additional cables are optional). Internal: Two internal USB 2.0 for embedded hypervisor.
Systems management	UEFI, Lenovo Integrated Management Module 2 (IMM2) with Renesas SH7757 controller, Predictive Failure Analysis, light path diagnostics panel, automatic server restart, and remote presence. Support for IBM Flex System Manager, IBM Systems Director, and Lenovo ServerGuide™.
Security features	Power-on password, administrator's password, and Trusted Platform Module V1.2.
Video	Matrox G200eR2 video core with 16 MB video memory that is integrated into the IMM2. Maximum resolution is 1600x1200 at 75 Hz with 16 M colors.
Limited warranty	Three-year customer-replaceable unit and onsite limited warranty with 9x5/NBD.
Operating systems supported	Microsoft Windows Server, Red Hat Enterprise , SUSE Linux Enterprise Server, VMware vSphere. For details, see 4.17, “Operating systems support” on page 70.

Components	Specification (per node, except where noted)
Service and support	Optional service upgrades are available through Lenovo ServicePac offerings: Four-hour or 2-hour response time, 24-hour or 8-hour fix time, 1-year or 2-year warranty extension, and remote technical support for Lenovo hardware and some Lenovo and OEM software.
Dimensions	Width: 435 mm (17.14 in.), height 55 mm (2.19 in.), depth 500 mm (19.7 in.)
Weight	Maximum weight: 12.25 kg (27 lbs).

Figure 4-3 shows the components on the system board of the x6 Compute Node.

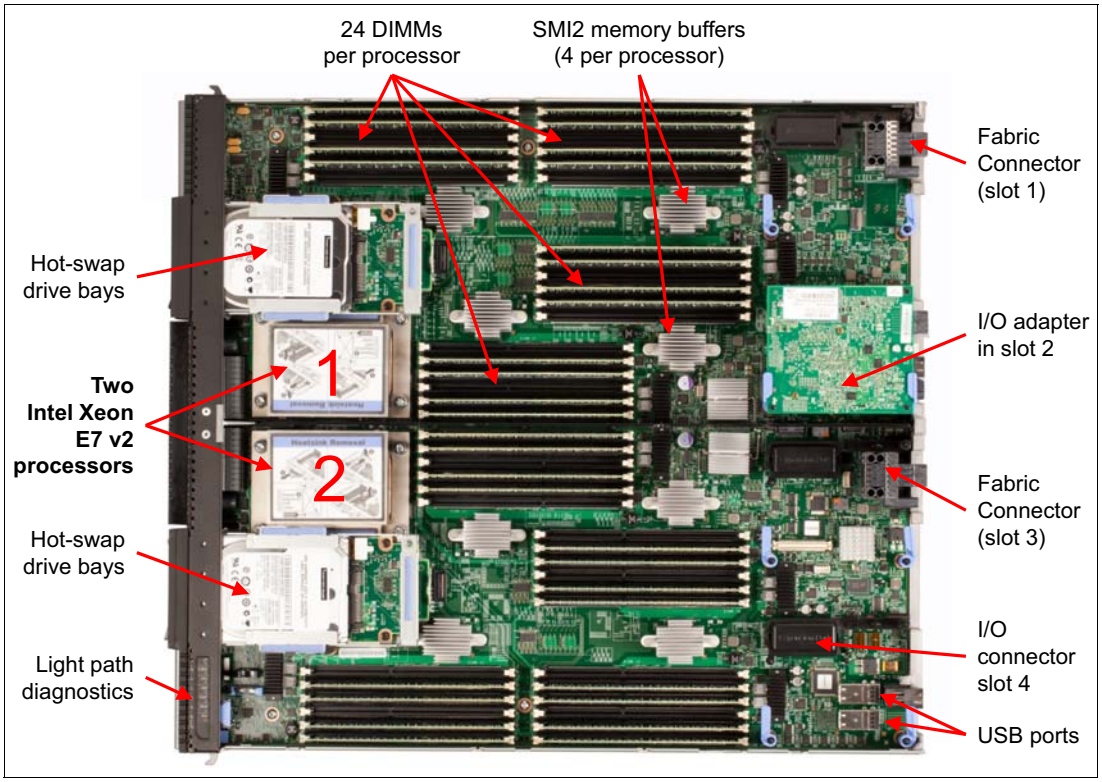

Figure 4-3 Layout of the X6 Compute Node system board

4.3 Models

The current Flex System X6 models are shown in Table 4-2.

Table 4-2 Standard models of the X6 Compute Node, type 7903

Model	Intel Processor[a]	Memory	RAID	Disk bays (used/max)	10GbE Embedded	I/O slots (used/max)
x280 X6 (2-socket only)						
7903-A2x	2x E7-2850 v2 12C 2.3GHz 24MB 1600MHz 105W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-B2x	2x E7-2880 v2 15C 2.5GHz 37.5MB 1600MHz 130W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4

Model	Intel Processor[a]	Memory	RAID	Disk bays (used/max)	10GbE Embedded	I/O slots (used/max)
7903-C2x	2x E7-2890 v2 15C 2.8GHz 37.5MB 1600MHz 155W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
x480 X6 (2-socket, can scale up to 4-socket)						
7903-D2x	2x E7-4809 v2 6C 1.9GHz 12MB 1333MHz 105W	2 x 16GB (1333MHz)[b]	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-F2x	2x E7-4820 v2 8C 2.0GHz 16MB 1600MHz 105W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-G2x	2x E7-4830 v2 10C 2.2GHz 20MB 1600MHz 105W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-H2x	2x E7-4850 v2 12C 2.3GHz 24MB 1600MHz 105W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-J2x	2x E7-4870 v2 15C 2.3GHz 30MB 1600MHz 130W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-L2x	2x E7-4880 v2 15C 2.5GHz 37.5MB 1600MHz 130W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
x880 X6 (2-socket, can scale up to 4-socket or 8-socket)						
7903-N2x	2x E7-8850 v2 12C 2.3GHz 24MB 1600MHz 105W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-Q2x	2x E7-8880 v2 15C 2.5GHz 37.5MB 1600MHz 130W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4
7903-R2x	2x E7-8890 v2 15C 2.8GHz 37.5MB 1600MHz 155W	2 x 16GB 1600 MHz	ServeRAID M1200e	2.5” HS (0 / 2) Open bay	No	0 / 4

a. Processor detail: Processor quantity and model, cores, core speed, L3 cache, memory speed, and power consumption.
b. Model D2x includes DIMMs that are rated at 1600 MHz but operate at 1333 MHz to match the processor memory speed.

4.4 Chassis support

The Lenovo Flex System X6 Compute Nodes are supported in the Lenovo Flex System Enterprise Chassis.

Up to seven X6 Compute Nodes can be installed in the chassis in 10U of rack space. The number of scaled systems that can be installed is as follows:

- Two-socket systems: Up to 7 systems in an Enterprise Chassis
- Four-socket systems: Up to 3 systems in an Enterprise Chassis
- Eight-socket systems: 1 system in an Enterprise Chassis

It might also be possible to populate the remaining bays within the chassis with other standard or double width nodes. The actual number of X6 systems that can be powered on in a chassis also depends on the following factors:

- The TDP power rating for the processors that are installed in the X6 node.
- The number of power supplies installed in the chassis.
- The capacity of the power supplies installed in the chassis (2100 W or 2500 W).
- The power redundancy policy used in the chassis (N+1 or N+N).

Table 2-2 on page 17 provides guidelines about what number of X6 systems can be powered on in the Lenovo Flex System Enterprise Chassis, based on the type and number of power supplies installed.

4.5 Scalability

The X6 Compute Nodes can scale to 4-socket or 8-socket, depending on the processor installed into the node.

- x880 X6 node with Intel Xeon E7-8800 v2 family processors, can scale up to 8-socket, and supports 2-socket, 4-socket, and 8-socket configurations. All processors in these 2-socket, 4-socket, and 8-socket configurations must be identical.
- x480 X6 node with Intel Xeon E7-4800 v2 family processors can scale up to 4-socket, and it supports 2-socket and 4-socket configurations. All processors in these 2-socket and 4-socket configurations must be identical.
- x280 X6 node with Intel Xeon E7-2800 v2 family processors supports only 2-socket configurations. The processors in 2-socket configurations must be identical.

The scaled X6 Compute Nodes are connected through a front interconnect system. The interconnection has a QPI bus plus the sideband signals necessary for proper operation.

There are two scalability connector assemblies for X6 nodes, as listed in Table 4-3.

Table 4-3 Scalability options for X6 Compute Nodes

Part number	Feature code	Description
00Y3871	A4D8	Lenovo Flex System x880 X6 4-Socket Scalability Connector (used to connect two x480 or x880 Compute Nodes together)
00Y3874	A4D9	Lenovo Flex System x880 X6 8-Socket Scalability Connector (used to connect four x880 Compute Nodes together)

The Lenovo Flex System x880 X6 4-Socket Scalability Connector shown at Figure 4-4.

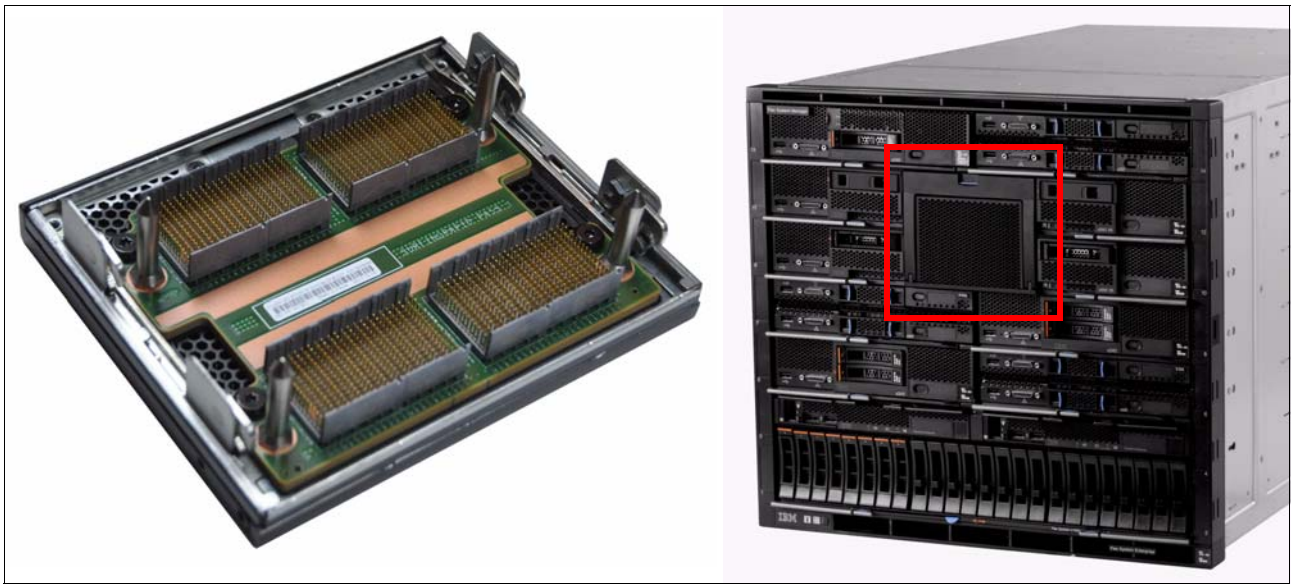

Figure 4-4 Lenovo Flex System x880 X6 4-Socket Scalability Connector

Lenovo Flex System x880 X6 4-Socket Scalability Connector allows for two x480 X6 or x880 X6 nodes to form a dual-node four-socket scaling complex. The scalability connection effectively connects all processors together through their QPI links. Figure 4-5 shows the block diagram for this connector.

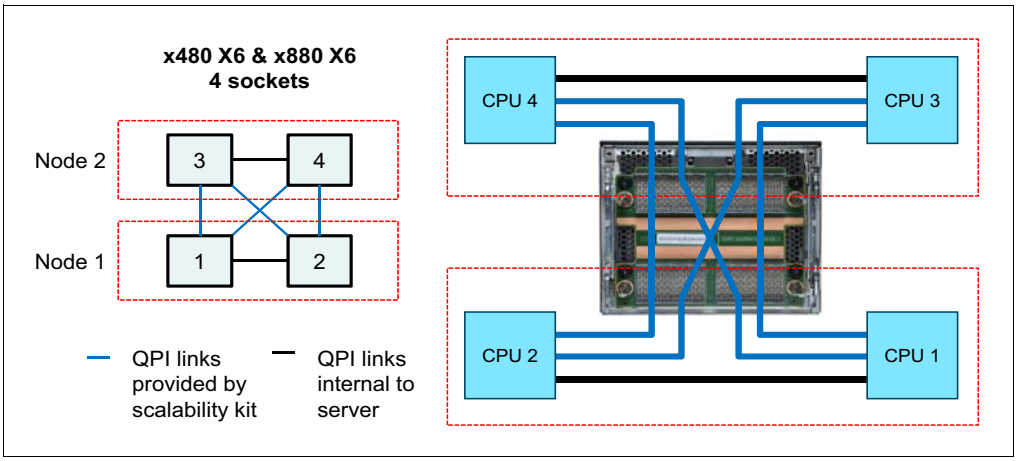

Figure 4-5 4-socket Scalability Connector: QPI connections

The Lenovo Flex System x880 X6 8-Socket Scalability Connector is shown in Figure 4-6 on page 53.

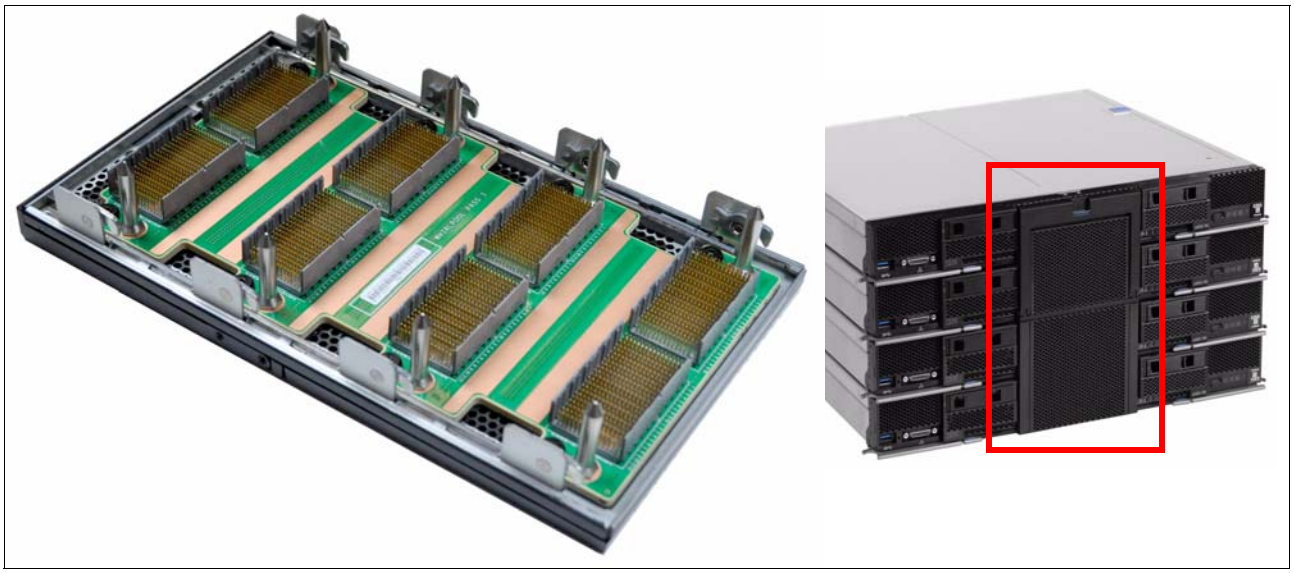

Figure 4-6 Lenovo Flex System x880 X6 8-Socket Scalability Connector

Lenovo Flex System x880 X6 8-Socket Scalability Connector allows for four x880 X6 nodes to form a 4-node, 8-socket scaling complex. Figure 4-7 shows the block diagram for this connector.

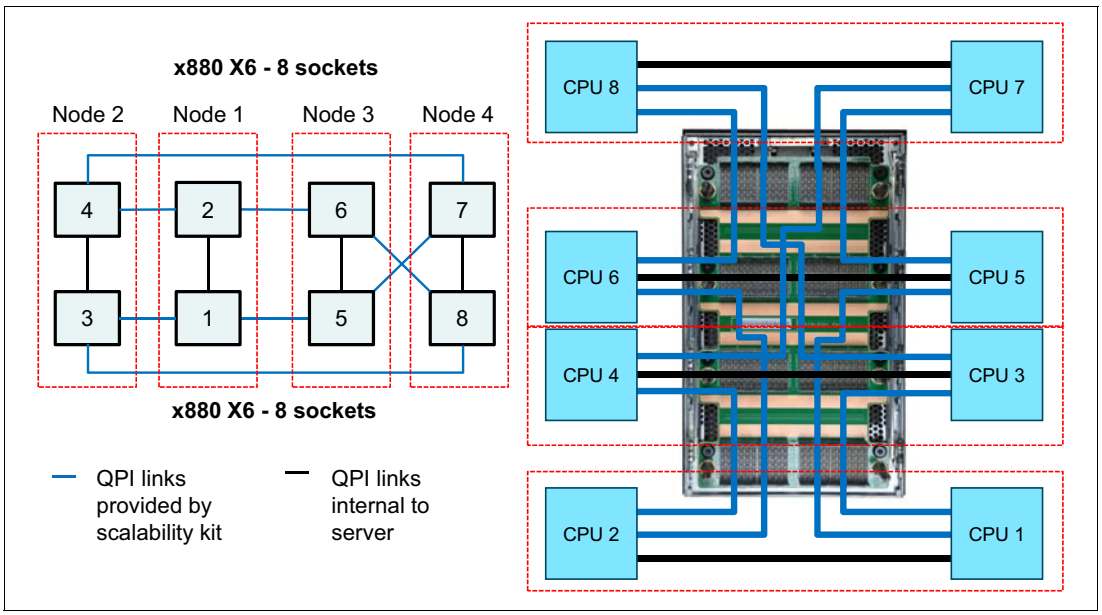

Figure 4-7 8-socket Scalability Connector diagram

The X6 Compute Nodes also offer three scalability LEDs on the front of the server to aid with problem determination. The primary scalability LED is near the power button. The other two LEDs are to the right of the Scalability Connector, as shown in Figure 4-8 on page 54.

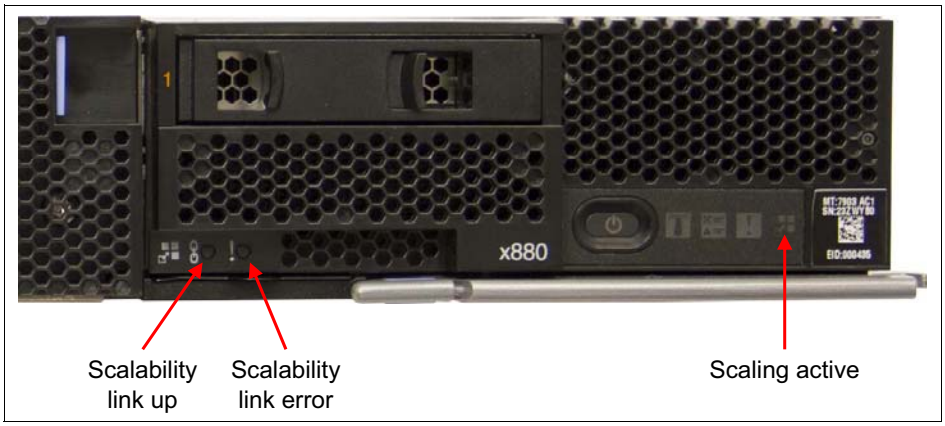

Figure 4-8 Scalability LEDs

The meaning of each of LEDs is described in the Table 4-4

Table 4-4 Scalability LEDs description.

LED	Color	Description
Link Up	Green	This LED lights when the system is configured as a multi-node system.
Link Error	Yellow	This LED lights when there is an error with the partitioning of the nodes.
Scaling	White	This LED lights on each node in a scalable complex when the nodes are operating as a scaled system. This LED does not light when the node is operating independently.

4.6 Processor options

The X6 Compute Node supports the processor options that are listed in Table 4-5. The server supports two Intel Xeon E7-2800 v2, E7-4800 v2, or E7-8800 v2 processors, depending on model. The table also shows which server models have each processor standard. If no corresponding model for a particular processor is listed, the processor is available only through the configure-to-order (CTO) process.

No part numbers: All Flex System X6 Compute Nodes have two processors as a standard. There is no option part number for processors, only feature codes.

Table 4-5 Supported processors for the X6 Compute Node

Feature code[a]	Intel Xeon processor description	Memory bus speed (RAS / Performance)[b]	Models where used
Lenovo Flex System x280 X6 Compute Node			
A4E0 / A4DF	Xeon E7-2850 v2 12C 2.3GHz 24 MB Cache 1600MHz 105W	1600 / 1066 MHz	-
A4E1 / A4DG	Xeon E7-2870 v2 15C 2.3GHz 30 MB Cache 1600MHz 130W	1600 / 1333 MHz	-
A4E2 / A4DE	Xeon E7-2880 v2 15C 2.5GHz 37.5 MB Cache 1600MHz 130W	1600 / 1333 MHz	-

Feature code[a]	Intel Xeon processor description	Memory bus speed (RAS / Performance)[b]	Models where used
A4E3 / A4DH	Xeon E7-2890 v2 15C 2.8GHz 37.5 MB Cache 1600MHz 155W	1600 / 1333 MHz	-
Lenovo Flex System x480 X6 Compute Node			
A4E4 / A4DJ	Xeon E7-4809 v2 6C 1.9GHz 12 MB Cache 1333MHz 105W	1333 / 1066 MHz	D2x
A4E5 / A4DK	Xeon E7-4820 v2 8C 2.0GHz 16 MB Cache 1600MHz 105W	1600 / 1066 MHz	-
A4E6 / A4DL	Xeon E7-4830 v2 10C 2.2GHz 20 MB Cache 1600MHz 105W	1600 / 1066 MHz	-
A4E7 / A4DM	XeonE7-4850 v2 12C 2.3GHz 24 MB Cache 1600MHz 105W	1600 / 1066 MHz	H2x
A4E8 / A4DN	Xeon E7-4860 v2 12C 2.6GHz 30 MB Cache 1600MHz 130W	1600 / 1333 MHz	-
A4E9 / A4DP	Xeon E7-4870 v2 15C 2.3GHz 30 MB Cache 1600MHz 130W	1600 / 1333 MHz	-
A4EA / A4DQ	Xeon E7-4880 v2 15C 2.5GHz 37.5 MB Cache 1600MHz 130W	1600 / 1333 MHz	L2x
A4EB / A4DR	Xeon E7-4890 v2 15C 2.8GHz 37.5 MB Cache 1600MHz 155W	1600 / 1333 MHz	-
Lenovo Flex System x880 X6 Compute Node			
A4EC / A4DS	Xeon E7-8850 v2 12C 2.3GHz 24 MB Cache 1600MHz 105W	1600 / 1066 MHz	-
A4EE / A4DU	Xeon E7-8870 v2 15C 2.3GHz 30 MB Cache 1600MHz 130W	1600 / 1333 MHz	-
A4EG / A4DW	Xeon E7-8880 v2 15C 2.5GHz 37.5 MB Cache 1600MHz 130W	1600 / 1333 MHz	-
A4EH / A4DX	Xeon E7-8890 v2 15C 2.8GHz 37.5 MB Cache 1600MHz 155W	1600 / 1333 MHz	-
A4EK / A4DZ	Xeon E7-8893 v2 6C 3.4GHz 37.5 MB Cache 1600MHz 155W	1600 / 1333 MHz	-
A4EJ / A4DY	Xeon E7-8891 v2 10C 3.2GHz 37.5 MB Cache 1600MHz 155W	1600 / 1333 MHz	-
A4ED / A4DT	Xeon E7-8857L v2 12C 3.0GHz 30 MB Cache 1600MHz 130W	1600 / 1333 MHz	-
A4EF / A4DV	Xeon E7-8880L v2 15C 2.2GHz 37.5 MB Cache 1600MHz 105W	1600 / 1333 MHz	-

a. The first feature code is for CPU 1, and the second feature code is for CPU 2.
b. The processors support two memory modes: RAS mode (Lockstep mode) and Performance mode (Independent mode). In Performance mode, the SMI2 link operates at twice the memory bus speed shown.

4.7 Memory options

As introduced in 3.6, “Memory architecture” on page 35, the X6 Compute Nodes support DDR3 memory operating at speeds up to 1600 MHz with 24 DIMMs per processor:

- A 2-socket configuration supports up to 48 DIMMs.
- A 4-socket scaled configuration supports up to 96 DIMMs.
- An 8-socket configuration supports up to 192 DIMMs.

Each processor has four memory channels to memory buffers implemented by using Scalable Memory Interface generation 2 (SMI2) chips. Each memory buffer has two memory channels and implements three DIMMs per channel.

The Intel Xeon E7-2800 v2, E7-4800 v2 and E7-8800 v2 processors support two memory modes: Performance mode and RAS (or *Lockstep*) mode. These are described in 3.6.1, “Operational modes” on page 36.

The following rules apply when selecting the memory configuration:

- The X6 node supports RDIMMs and LRDIMMs.
- LRDIMMs and RDIMMs cannot be mixed within a compute node or a scaled complex.
- Mixing 1.5 V and 1.35 V DIMMs in the same server is supported. In such a case, all DIMMs operate at 1.5 V.
- In RAS (Lockstep) mode, DIMMs must be installed in a pair, and the SMI link operates at the speed of the memory bus.
- The maximum number of ranks that is supported per channel is eight (except for Load Reduced DIMMs, where more than eight ranks are supported, because one quad-rank LRDIMM provides the same electrical load on a memory bus as a single-rank RDIMM, and one 8-rank LRDIMM provides the same electrical load on a memory bus as a dual-rank RDIMM).
- All DIMMs in all processor memory channels operate at the same speed, which is determined as the lowest value of the following components:
  - Memory speed that is supported by a specific processor.
  - Lowest of maximum operating speeds for selected memory configuration, depending on rated speed, operating voltage, and quantity of DIMMs per channel, as shown under the “Maximum operating speed” section in Table 4-7 on page 57.

Table 4-6 lists the memory options that are available for the X6 server. DIMMs can be installed one at a time, but for performance reasons, install them in sets of four (one for each of the memory channels). A total of 48 DIMMs is the maximum number supported.

Table 4-6 Memory options for the X6 node

Part number	Feature code	Description	Maximum supported	Models where used
Registered DIMM (RDIMM) modules				
00D5024	A3QE	4GB (1x4GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	48	-
00D5036	A3QH	8GB (1x8GB, 1Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	48	-
46W0672	A3QM	16GB (1x16GB, 2Rx4, 1.35V) PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	48	All models
Load Reduced DIMM (LRDIMM) modules				
46W0676	A3SR	32GB (1x32GB, 4Rx4, 1.35V)PC3L-12800 CL11 ECC DDR3 1600MHz LP LRDIMM	48	-
46W0741	A451	64GB (1x64GB, 8Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz LP LRDIMM	48[a]	-

a. 64 GB LRDIMMs are supported only in 2-socket or 4-socket configurations.

Table 4-7 on page 57 shows the characteristics of the supported DIMMs and the memory speeds. The cells highlighted in gray indicate that the X6 servers support higher memory frequencies or larger memory capacity, or both, than the Intel processor specification defines.

Memory speed: In Performance mode, memory channels operate independently, and the SMI2 link operates at twice the DDR3 speed. In RAS mode, two channels operate synchronously, and the SMI2 link operates at the DDR3 speed.

Table 4-7 Maximum memory speeds

DIMM specification	RDIMM				LR-DIMM			
Ranks	Single-rank DIMM		Dual-rank DIMM		Quad-rank LR-DIMM		8-rank LR-DIMM	
Part numbers	00D5024 (4 GB) 00D5036 (8 GB)		46W0672 (16 GB)		46W0676 (32 GB)		46W0741 (64 GB)	
Rated speed	1600 MHz		1600 MHz		1600 MHz		1333 MHz	
Rated voltage	1.35 V		1.35 V		1.35 V		1.35 V	
Operating voltage	1.35 V	1.5 V	1.35 V	1.5 V	1.35 V	1.5 V	1.35 V	1.5 V
Max. DIMM capacity	8 GB	8 GB	16 GB	16 GB	32 GB	32 GB	64 GB	64 GB
Max. memory capacity (2 socket)	384 GB	384 GB	0.75 TB	0.75 TB	1.5 TB	1.5 TB	3 TB	3 TB
Max. memory capacity (4 socket)	0.75 TB	0.75 TB	1.5 TB	1.5 TB	3 TB	3 TB	6 TB	6 TB
Max. memory capacity (8 socket)	1.5 TB	1.5 TB	3 TB	3 TB	6 TB	6 TB	No support[a]	No support[a]
Maximum operating speed: Performance mode (2:1 mode - SMI2 link operates at *twice* the DDR3 speed shown)								
1 DIMM per channel	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz
2 DIMMs per channel	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz
3 DIMMs per channel	1066 MHz	1333 MHz	1066 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz
Maximum operating speed: RAS mode (1:1 mode - SMI2 link operates at the DDR3 speed shown)								
1 DIMM per channel	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz
2 DIMMs per channel	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1600 MHz	1333 MHz	1333 MHz
3 DIMMs per channel	1066 MHz	1333 MHz	1066 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz	1333 MHz

a. 64 GB DIMMs are supported only in 2-socket or 4-socket configurations.

4.8 Memory population order

The X6 Compute Node has two main memory operation modes: Performance mode and RAS mode, with two additional resiliency modes: Mirroring and Rank Sparing.

Performance or RAS modes can be complimented with one of the additional resiliency modes, so overall there can be six operation modes:

- Performance mode (also referred to as *Independent* mode)
- RAS mode (also referred to as *Lockstep* mode)

- Performance mode with Mirroring
- Performance mode with Rank Sparing
- RAS mode with Mirroring
- RAS mode with Rank Sparing

For more information about memory operation modes, see 3.6.1, “Operational modes” on page 36.

The diagram in Figure 4-9 shows DIMM numbering on the system board.

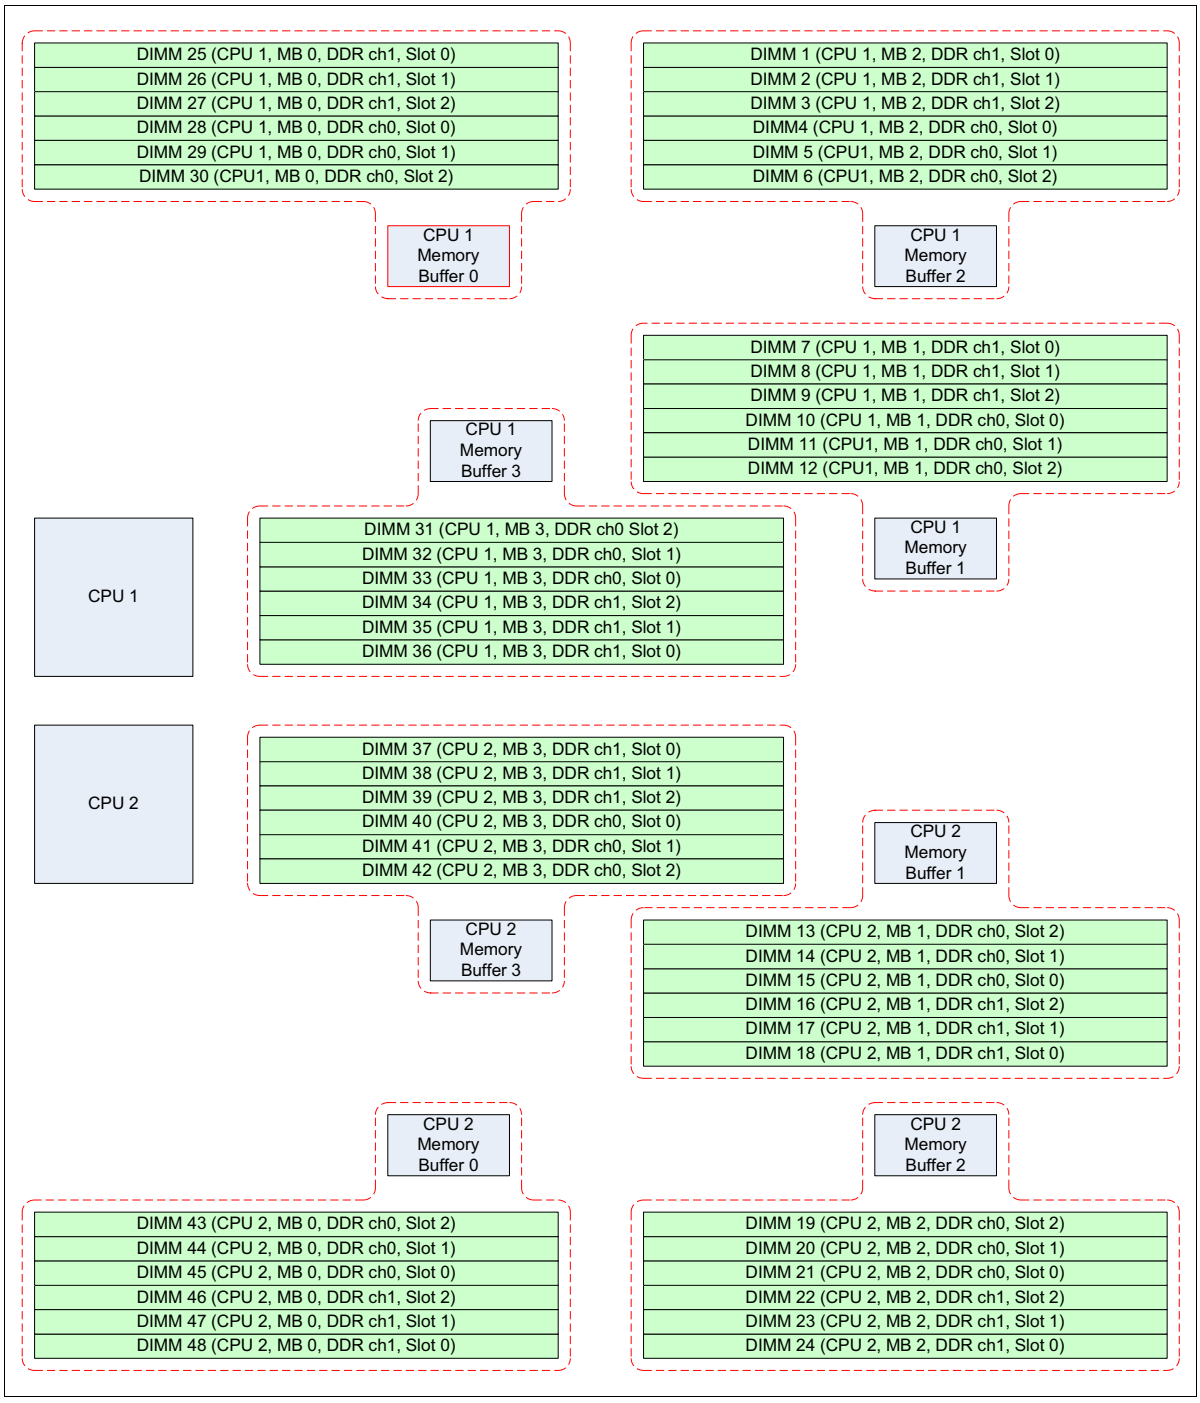

Figure 4-9 DIMM numbering on the x6 system board

DIMM slot locations: For CPU 1, DIMM slot locations for DIMMs 31 - 36 are inverse to all other DIMMs for CPU1. CPU2 DIMM slot locations for DIMMs 37 - 42 are inverse to all other DIMMs for CPU2.

In all cases, the numbering is in relation to the connected memory buffer, as indicated by the dotted red outlines. For example, channel 1 slot 0 is always farthest away from the memory buffer, regardless of the orientation of the DIMM numbering.

Depending on the selected operation mode, use appropriate DIMMs in a specific installation order. Table 4-8 shows DIMM placement order depending upon operation mode.

The DIMM installation order for Rank Sparing resiliency follows the main operation mode that it complements (Performance mode or RAS mode).

Table 4-8 DIMM placement order

DIMM order	Performance mode	Performance mode with memory mirroring	RAS mode	RAS mode with memory mirroring
1	DIMM 28	Not applicable	Not applicable	Not applicable
2	DIMM 45	DIMMs 4 and 28	DIMMs 25 and 28	Not applicable
3	DIMM 10	Not applicable	Not applicable	Not applicable
4	DIMM 15	DIMMs 21 and 45	DIMMs 45 and 48	DIMMs 1, 4, 25, and 28
5	DIMM 4	Not applicable	Not applicable	Not applicable
6	DIMM 21	DIMMs 10 and 33	DIMMs 7 and 10	Not applicable
7	DIMM 33	Not applicable	Not applicable	Not applicable
8	DIMM 40	DIMMs 15 and 40	DIMMs 15 and 18	DIMMs 21, 24, 45, and 48
9	DIMM 25	Not applicable	Not applicable	Not applicable
10	DIMM 48	DIMMs 1 and 25	DIMMs 1 and 4	Not applicable
11	DIMM 7	Not applicable	Not applicable	Not applicable
12	DIMM 18	DIMMs 24 and 48	DIMMs 21 and 24	DIMMs 7, 10, 33, and 36
13	DIMM 1	Not applicable	Not applicable	Not applicable
14	DIMM 24	DIMMs 7 and 36	DIMMs 33 and 36	Not applicable
15	DIMM 36	Not applicable	Not applicable	Not applicable
16	DIMM 37	DIMMs 18 and 37	DIMMs 37 and 40	DIMMs 15, 18, 37, and 40
17	DIMM 29	Not applicable	Not applicable	Not applicable
18	DIMM 44	DIMMs 5 and 29	DIMMs 26 and 29	Not applicable
19	DIMM 11	Not applicable	Not applicable	Not applicable
20	DIMM 14	DIMMs 20 and 44	DIMMs 44 and 47	DIMMs 2, 5, 26, and 29
21	DIMM 5	Not applicable	Not applicable	Not applicable
22	DIMM 20	DIMMs 11 and 32	DIMMs 8 and 11	Not applicable
23	DIMM 32	Not applicable	Not applicable	Not applicable

DIMM order	Performance mode	Performance mode with memory mirroring	RAS mode	RAS mode with memory mirroring
24	DIMM 41	DIMMs 14 and 41	DIMMs 14 and 17	DIMMs 20, 23, 44 and 47
25	DIMM 26	Not applicable	Not applicable	Not applicable
26	DIMM 47	DIMMs 2 and 26	DIMMs 2 and 5	Not applicable
27	DIMM 8	Not applicable	Not applicable	Not applicable
28	DIMM 17	DIMMs 23 and 47	DIMMs 20 and 23	DIMMs 8, 11, 32, and 35
29	DIMM 2	Not applicable	Not applicable	Not applicable
30	DIMM 23	DIMMs 8 and 35	DIMMs 32 and 35	Not applicable
31	DIMM 35	Not applicable	Not applicable	Not applicable
32	DIMM 38	DIMMs 17 and 38	DIMMs 38 and 41	DIMMs 14, 17, 38, and 41
33	DIMM 30	Not applicable	Not applicable	Not applicable
34	DIMM 43	DIMMs 6 and 30	DIMMs 27 and 30	Not applicable
35	DIMM 12	Not applicable	Not applicable	Not applicable
36	DIMM 13	DIMMs 19 and 43	DIMMs 43 and 46	DIMMs 3, 6, 27, and 30
37	DIMM 6	Not applicable	Not applicable	Not applicable
38	DIMM 19	DIMMs 12 and 31	DIMMs 9 and 12	Not applicable
39	DIMM 31	Not applicable	Not applicable	Not applicable
40	DIMM 42	DIMMs 13 and 42	DIMMs 13 and 16	DIMMs 19, 22, 43, and 46
41	DIMM 27	Not applicable	Not applicable	Not applicable
42	DIMM 46	DIMMs 3 and 27	DIMMs 3 and 6	Not applicable
43	DIMM 9	Not applicable	Not applicable	Not applicable
44	DIMM 16	DIMMs 22 and 46	DIMMs 19 and 22	DIMMs 9, 12, 31, and 34
45	DIMM 3	Not applicable	Not applicable	Not applicable
46	DIMM 22	DIMMs 9 and 34	DIMMs 31 and 34	Not applicable
47	DIMM 34	Not applicable	Not applicable	Not applicable
48	DIMM 39	DIMMs 16 and 39	DIMMs 39 and 42	DIMMs 13, 16, 39, and 42

4.9 Internal disk controller

The X6 Compute Node has two 2.5-inch hot-swap drive bays that are accessible from the front of the server, as shown in Figure 4-2 on page 47. These bays are connected to the integrated 4-port LSI SAS 3004 12 Gbps SAS/SATA RAID-on-Chip (ROC) controller.

The integrated LSI SAS 3004 ROC includes the following features:

- Four-port controller with 12 Gbps throughput per port
- PCIe x4 Gen 2 host interface
- Two SAS ports that are routed internally to the two hot-swap drive bays
- Support for RAID levels 0, 1, 10, and 1E; optional RAID 5 and RAID 50

RAID 5 and RAID 50 are supported by a Lenovo Features on Demand (FoD) upgrade, as shown in Table 4-9.

Table 4-9 RAID 5 upgrade

Part number	Feature code	Description
00AE930	A5H5	ServeRAID M1200 Zero Cache/RAID 5 upgrade for Lenovo Systems (FoD)

4.10 Internal disk options

The 2.5-inch drive bays support SAS or SATA HDDs or SATA SSDs.

Table 4-10 lists the supported 2.5-inch drive options.

Table 4-10 2.5-inch internal disk options

Part number	Feature code		Description	Maximum supported
10K SAS hard disk drives				
00AJ146	A4TP		1.2TB 10K 6Gbps SAS 2.5" G3HS HDD	2
00AJ071	A4TN		900GB 10K 6Gbps SAS 2.5" SFF G3HS HDD	2
00AJ091	A4TM		600GB 10K 6Gbps SAS 2.5" SFF G3HS HDD	2
00AJ096	A4TL		300GB 10K 6Gbps SAS 2.5" SFF G3HS HDD	2
10K and 15K SAS self-encrypting drives (SEDs)[a]				
00AJ116	A4U2		146GB 15K 6Gbps SAS 2.5" G3HS SED	2
00AJ151	A4U1		1.2TB 10K 6Gbps SAS 2.5" G3HS SED	2
00AJ076	A4U0		900GB 10K 6Gbps SAS 2.5" SFF G3HS SED	2
00AJ101	A4TZ		600GB 10K 6Gbps SAS 2.5" SFF G3HS SED	2
00AJ106	A4TY		300GB 10K 6Gbps SAS 2.5" SFF G3HS SED	2
15K SAS hard disk drives				
00AJ126	A4TS		Lenovo 600GB 15K 6Gbps SAS 2.5" G3HS HDD	2
00AJ081	A4TR		Lenovo 300GB 15K 6Gbps SAS 2.5" G3HS HDD	2

Part number	Feature code		Description	Maximum supported
00AJ111	A4TQ		146GB 15K 6Gbps SAS 2.5" G3HS HDD	2
NL SATA drives				
00AJ141	A4TX		1TB 7.2K 6Gbps NL SATA 2.5" G3HS HDD	2
00AJ136	A4TW		500GB 7.2K 6Gbps NL SATA 2.5" G3HS HDD	2
00AJ131	A4TV		250GB 7.2K 6Gbps NL SATA 2.5" G3HS HDD	2
NL SAS drives				
00AJ086	A4TU		1TB 7.2K 6Gbps NL SAS 2.5" G3HS HDD	2
00AJ121	A4TT		500GB 7.2K 6Gbps NL SAS 2.5" G3HS HDD	2
Enterprise SSDs				
00AJ222	A4UD		1.6TB SAS 2.5" MLC G3HS Enterprise SSD	2
00AJ217	A4UC		800GB SAS 2.5" MLC G3HS Enterprise SSD	2
00AJ212	A4UB		400GB SAS 2.5" MLC G3HS Enterprise SSD	2
00AJ207	A4UA		200GB SAS 2.5" MLC G3HS Enterprise SSD	2
00AJ166	A4U5		S3700 800GB SATA 2.5" MLC G3HS Enterprise SSD	2
00AJ161	A4U4		S3700 400GB SATA 2.5" MLC G3HS Enterprise SSD	2
00AJ156	A4U3		S3700 200GB SATA 2.5" MLC G3HS Enterprise SSD	2
Enterprise Value SSDs				
00AJ395	A577		120GB SATA 2.5" MLC G3HS Enterprise Value SSD	2
00AJ400	A578		240GB SATA 2.5" MLC G3HS Enterprise Value SSD	2
00AJ405	A579		480GB SATA 2.5" MLC G3HS Enterprise Value SSD	2
00AJ410	A57A		800GB SATA 2.5" MLC G3HS Enterprise Value SSD	2

a. Supports self-encrypting drive (SED) technology. For more information, see the Self-Encrypting Drives for Lenovo System x web page: http://lenovopress.com/tips0761

4.10.1 Support for 1.8-inch SSDs

The X6 Compute Nodes currently do not support 1.8-inch solid-state drives.

4.11 I/O expansion options

The X6 node has four I/O expansion connectors for attaching I/O adapters, as shown in Table 4-10 on page 63.

Installing I/O adapters allows the server to connect to switch modules in the Lenovo Flex System Enterprise Chassis. Table 4-10 shows the location of the four I/O expansion slots.

Dual-ASIC adapters: Slots 3 and 4 support only a subset of the adapters that are supported in slots 1 and 2. Slots 3 and 4 do not support dual-ASIC adapters due to the PCIe lanes that are routed to slots 3 and 4.

See 4.12, “Network adapters” on page 65 and 4.13, “Storage host bus adapters” on page 66 for specifics.

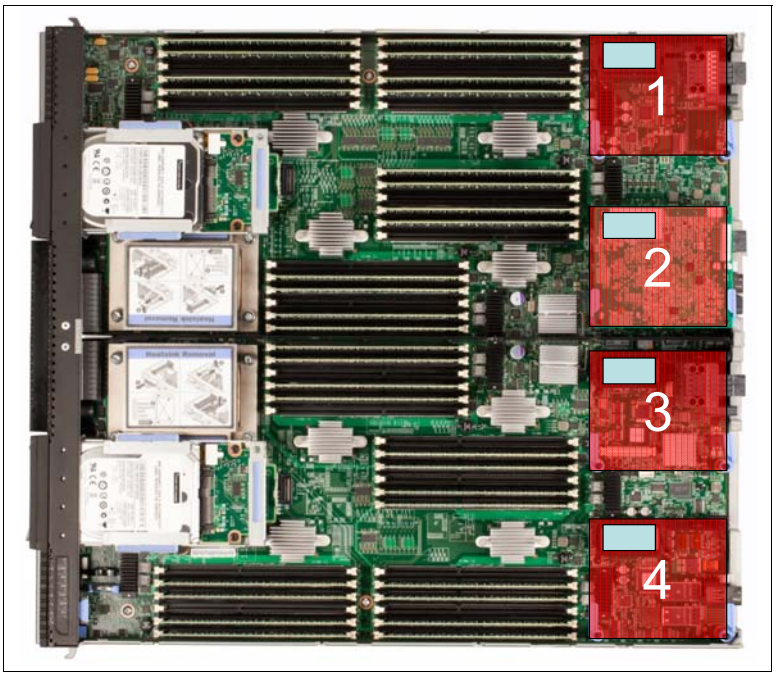

Figure 4-10 Location of the I/O adapters in the Lenovo Flex System X6 Compute Node

A compatible switch or pass-through module must be installed in the corresponding I/O bays in the chassis, as indicated in Table 4-11. Installing two switches means that all ports of the adapter are enabled, which improves performance and network availability.

Table 4-11 Adapter to I/O bay correspondence

I/O adapter slot in the X6 Compute Nodes	Port on the adapter	Corresponding I/O module bay in the chassis
Slot 1	Port 1	Module bay 1
	Port 2	Module bay 2
	Port 3 (for 4 and 8-port cards)	Module bay 1
	Port 4 (for 4 and 8-port cards)	Module bay 2
	Port 5 (for 8-port cards)	Module bay 1
	Port 6 (for 8-port cards)	Module bay 2
	Port 7 (for 8-port cards)	Module bay 1
	Port 8 (for 8-port cards)	Module bay 2

I/O adapter slot in the X6 Compute Nodes	Port on the adapter	Corresponding I/O module bay in the chassis
Slot 2	Port 1	Module bay 3
	Port 2	Module bay 4
	Port 3 (for 4 and 8-port cards)	Module bay 3
	Port 4 (for 4 and 8-port cards)	Module bay 4
	Port 5 (for 8-port cards)	Module bay 3
	Port 6 (for 8-port cards)	Module bay 4
	Port 7 (for 8-port cards)	Module bay 3
	Port 8 (for 8-port cards)	Module bay 4
Slot 3 (Dual-ASIC adapters not supported)	Port 1	Module bay 1
	Port 2	Module bay 2
Slot 4 (Dual-ASIC adapters not supported)	Port 1	Module bay 3
	Port 2	Module bay 4

Figure 4-11 shows the location of the switch bays in the rear of the Enterprise Chassis.

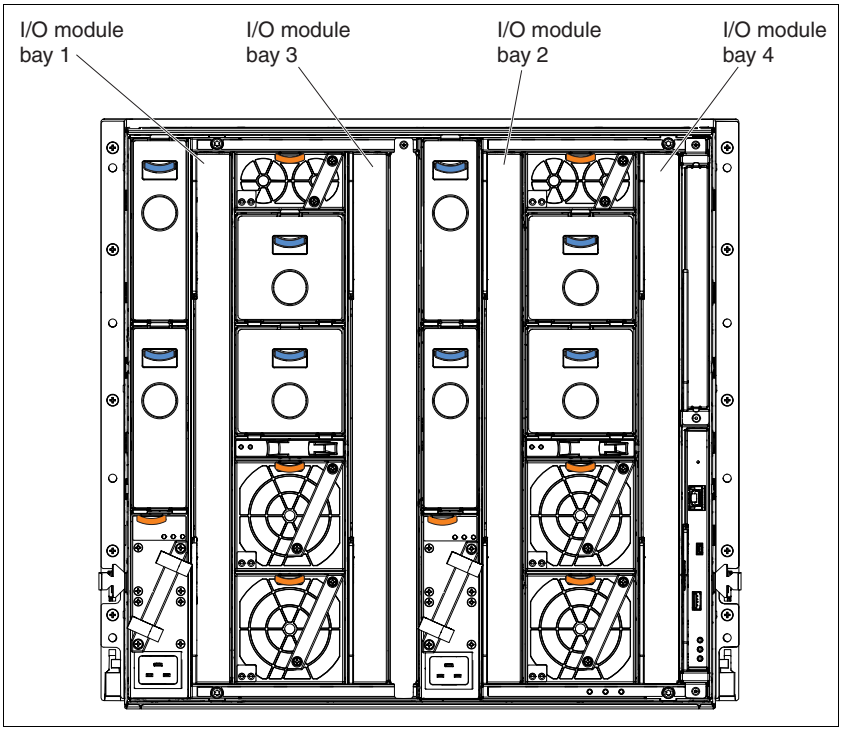

Figure 4-11 Locations of the I/O modules

Figure 4-12 shows how the two port adapters are connected to switches that are installed in the I/O module bays in an Enterprise Chassis.

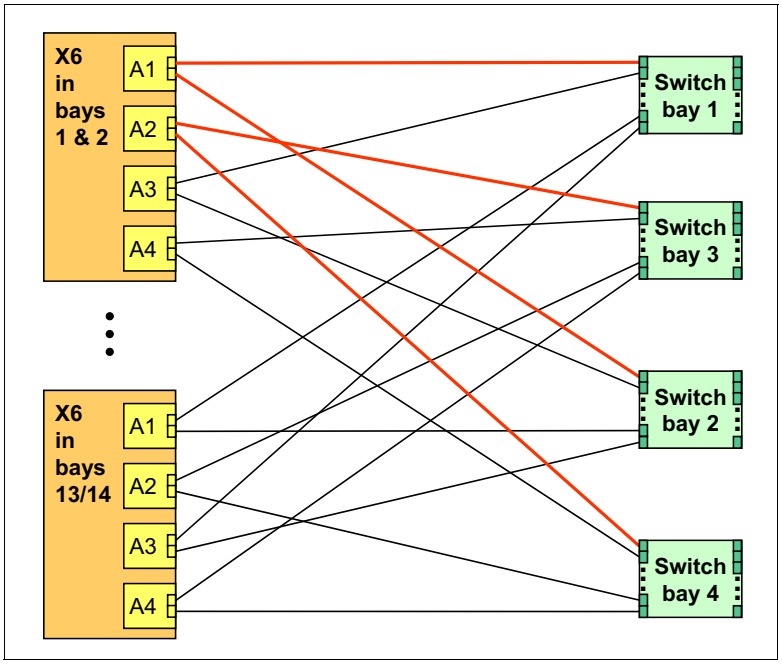

Figure 4-12 Logical layout of the interconnects between I/O adapters and I/O modules

4.12 Network adapters

A network adapter is required to be installed in I/O slot 1 (and, optionally, slot 2 for redundancy and additional ports).

Table 4-12 lists the supported network adapters and upgrades. Adapters can be installed in the slot as indicated in the table. Compatible switches must be installed in the corresponding bays of the chassis.

Note: The following adapters with dual ASICs are not supported in slots 3 and 4:

- Lenovo Flex System CN4054R 10Gb Virtual Fabric Adapter
- Lenovo Flex System EN2024 4-port 1Gb Ethernet Adapter

Table 4-12 Network adapters

Part number	Feature code	Description	Number of ports	Maximum supported[a]	Slots supported
40Gb Ethernet					
90Y3482	A3HK	Lenovo Flex System EN6132 2-port 40Gb Ethernet Adapter	2	4	1, 2, 3, 4
10Gb Ethernet					
00Y3306	A4K2	Lenovo Flex System CN4054R 10Gb Virtual Fabric Adapter	4	2	1, 2[b]

Part number	Feature code	Description	Number of ports	Maximum supported[a]	Slots supported
90Y3558	A1R0	Lenovo Flex System CN4054 Virtual Fabric Upgrade (FoD) FCoE/iSCSI upgrade for 00Y3306; one per adapter	License	2[c]	Not applicable
88Y5920	A4K3	Lenovo Flex System CN4022 2-port 10Gb Converged Adapter	2	4	1, 2, 3, 4
90Y3466	A1QY	Lenovo Flex System EN4132 2-port 10Gb Ethernet Adapter	2	4	1, 2, 3, 4
1Gb Ethernet					
49Y7900	A10Y	Lenovo Flex System EN2024 4-port 1Gb Ethernet Adapter	4	2	1, 2[b]
InfiniBand					
90Y3454	A1QZ	Lenovo Flex System IB6132 2-port FDR InfiniBand Adapter	2	2	2, 4

a. For X6 models with Embedded 10Gb Virtual Fabric controller as standard, the Compute Node Fabric Connectors occupy the same space as the I/O adapters in I/O slots 1 and 3. Adapters are not supported in those slots.
b. Adapters with two ASICs are not supported in slots 3 and 4 due to the available PCIe lanes in those slots.
c. One Upgrade License per CN4054R adapter is required.

4.13 Storage host bus adapters

Table 4-13 lists storage host bus adapters (HBAs) that are supported by the X6 node.

Note: The following adapter with two ASICs is not supported in slots 3 and 4:

- Lenovo Flex System FC5054 4-port 16Gb FC adapter

Table 4-13 Storage adapters

Part number	Feature code	Description	Number of ports	Maximum supported per 2-socket node	Slots supported
69Y1942	A1BQ	Lenovo Flex System FC5172 2-port 16Gb FC adapter	2	2	2, 4
95Y2391	A45S	Lenovo Flex System FC5054 4-port 16Gb FC adapter	4	1	2
95Y2386	A45R	Lenovo Flex System FC5052 2-port 16Gb FC adapter	2	2	2, 4
88Y6370	A1BP	Lenovo Flex System FC5022 2-port 16Gb FC adapter	2	2	2, 4
69Y1938	A1BM	Lenovo Flex System FC3172 2-port 8Gb FC adapter	2	2	2, 4
95Y2375	A2N5	Lenovo Flex System FC3052 2-port 8Gb FC adapter	2	2	2, 4

4.14 Integrated virtualization

The X6 Compute Node offers USB flash drive options that are preinstalled with versions of VMware ESXi. This software is an embedded version of VMware ESXi and is contained on the flash drive, without requiring any disk space. The USB memory key plugs into one of the two internal USB ports on the system board. The supported USB memory keys are listed in Table 4-14.

Table 4-14 Virtualization options

Part number	Feature code	Description	Maximum supported
41Y8300	A2VC	Lenovo USB Memory Key for VMware ESXi 5.0	1
41Y8307	A383	Lenovo USB Memory Key for VMware ESXi 5.0 Update 1	1
41Y8311	A2R3	Lenovo USB Memory Key for VMware ESXi 5.1	1
41Y8382	A4WZ	Lenovo USB Memory Key for VMware ESXi 5.1Update 1	1
41Y8385	A584	Lenovo USB Memory Key for VMware ESXi 5.5	1
41Y8298	A2G0	Lenovo Blank USB Memory Key for VMware ESXi Downloads[a]	2

a. The blank USB Memory Key requires the download of the VMware vSphere (ESXi) Hypervisor with Lenovo Customization image, which is available on the Lenovo x86 solutions for VMware web page: http://ibm.com/systems/x/os/vmware/

There are two types of USB keys: preload keys or blank keys. Blank keys allow you to download a Lenovo customized version of ESXi and load it onto the key. Each server supports installing one or two keys but only the following combinations:

- One preload key (keys that are preloaded at the factory)

- One blank key (a key for which you download the customized image)
- One preload key and one blank key
- Two blank keys

Two preload keys is an unsupported combination. Installing two preload keys prevents ESXi from booting. This is similar to the error described on the VMware Knowledge Base website, “ESXi host displays an error and fails to boot when two ESXi installations are on the same system (1035107):”

http://kb.vmware.com/kb/1035107

Having two keys that are installed provides a backup boot device. Both devices are listed in the boot menu, which allows you to boot from either device or to set one as a backup in case the first one becomes corrupted

4.15 Light path diagnostics panel

For quick problem determination when you are physically at the server, the X6 offers the following three-step guided path:

1. The Fault LED on the front panel
2. The light path diagnostics panel
3. LEDs next to key components on the system board

The X6 light path diagnostics panel is visible when you remove the server from the chassis. The panel is at the upper-right side of the compute node, as Figure 4-13 shows.

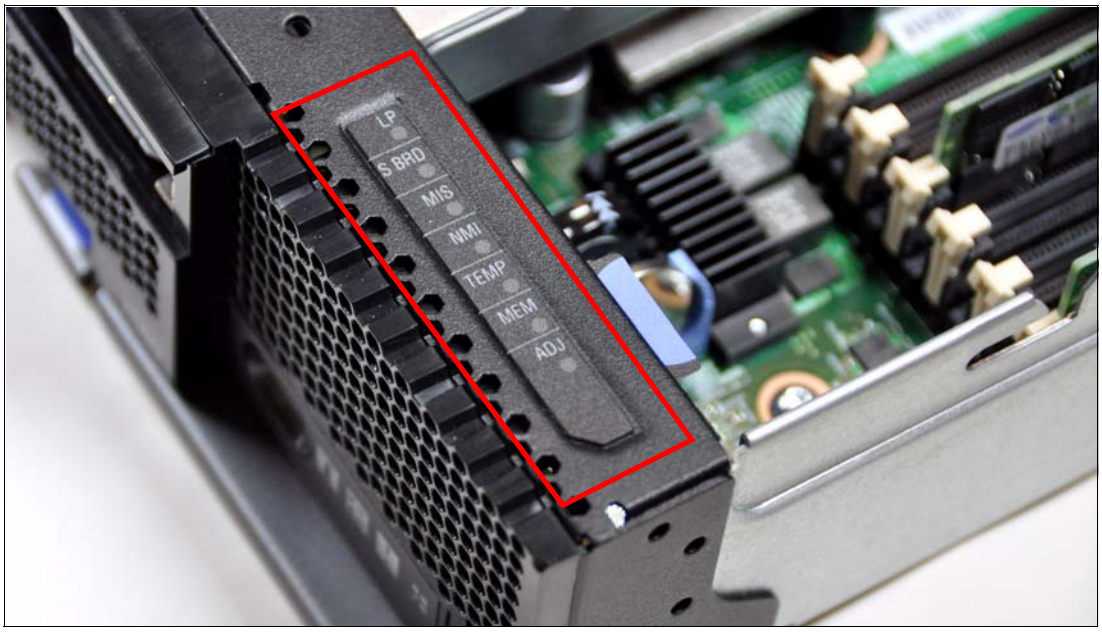

Figure 4-13 Location of X6 light path diagnostics panel

To illuminate the light path diagnostics LEDs, power off the compute node, slide it out of the chassis, and press the power button. The power button doubles as the light path diagnostics remind button when the server is removed from the chassis.

The meanings of the LEDs in the light path diagnostics panel are listed in Table 4-15.

Table 4-15 Light path diagnostic panel LEDs

LED	Meaning
LP	The light path diagnostics panel is operational.
S BRD	A system board error is detected.
MIS	A mismatch occurred between the processors, DIMMs, or HDDs within the configuration reported by POST.
NMI	A non-maskable interrupt (NMI) occurred.
TEMP	An over-temperature condition occurred that was critical enough to shut down the server.
MEM	A memory fault occurred. The corresponding DIMM error LEDs on the system board are also lit.
ADJ	A fault is detected in the adjacent expansion unit (if installed).

The server also includes LEDs for scalability. Those are described in 4.5, “Scalability” on page 51.

4.16 Remote management

The server contains a Lenovo Integrated Management Module II (IMMv2), which interfaces with the advanced management module in the chassis. The combination of these two components provides advanced service-processor control, monitoring, and an alerting function.

If an environmental condition exceeds a threshold or a system component fails, LEDs lit on the system board aid in diagnosis, the error is recorded in the event log, and you are alerted to the problem. A virtual presence capability is standard for remote server management.

Remote server management is provided through the following industry-standard interfaces:

- Intelligent Platform Management Interface (IPMI) Version 2.0
- Simple Network Management Protocol (SNMP) Version 3
- Common Information Model (CIM)
- Web browser

The server also supports virtual media and remote control features, which provide the following functions:

- Remotely view video with graphics resolutions up to 1600 x 1200 at 75 Hz with up to 23 bits per pixel, regardless of the system state
- Remotely access the server by using the keyboard and mouse from a remote client
- Map the CD or DVD drive, diskette drive, and USB flash drive on a remote client, and mapping ISO and diskette image files as virtual drives that are available for use by the server
- Upload a diskette image to the IMM2 memory and map it to the server as a virtual drive
- Capture blue-screen errors

4.17 Operating systems support

The X6 node supports the following operating systems:

- Microsoft Windows Server 2008 R2 Datacenter SP1
- Microsoft Windows Server 2008 R2 Enterprise SP1
- Microsoft Windows Server 2008 R2 Standard SP1
- Microsoft Windows Server 2012 Datacenter
- Microsoft Windows Server 2012 Enterprise
- Microsoft Windows Server 2012 Standard
- Microsoft Windows Storage Server 2012 Standard
- Microsoft Windows Server 2012 R2 Datacenter
- Microsoft Windows Server 2012 R2 Enterprise
- Microsoft Windows Server 2012 R2 Standard
- Microsoft Windows Storage Server 2012 R2 Standard
- Red Hat Enterprise Linux 6 Server x64 Edition U5
- Red Hat Enterprise Linux 7
- SUSE Linux Enterprise Server 11 for AMD64/EM64T SP3
- SUSE Linux Enterprise Server 11 with Xen for AMD64/EM64T SP3
- VMware vSphere 5.0
- VMware vSphere 5.0 Update 1
- VMware vSphere 5.1 U2
- VMware vSphere 5.1 Update 1 U2
- VMware vSphere 5.5

Support by some of these operating system versions is after the date of initial availability. Check the Lenovo ServerProven® web page titled “System x and xSeries® OS Support Information” for the latest information about the specific versions and service levels that are supported and any other prerequisites:

http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/nos/matrix.shtml

5

Hardware preparation

This chapter provides assistance for making configuration, monitoring, and maintenance decisions when implementing a Lenovo Flex System X6 Compute Node complex. It covers the following topics:

5.1 Integrated Management Module II (IMM2)

The Integrated Management Module II (IMM2) is the next generation of the IMM service processor, which provides advanced control and monitoring features to manage the Lenovo Flex System X6 Compute Node family. IMM2 enables easy console redirection with text, graphics, and full KVM support over the systems management connection through the Chassis Management Module (CMM).

These features allow you to display server activities from power-on to full remote operation, with interaction at nearly any time.

For more detailed information, see the *Integrated Management Module II User's Guide* available from:

`http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086346`

This section describes the initial setup tasks to perform after receiving your server. It includes the following topics:

5.1.1 IMM2 Overview

The IMM2 consolidates the service processor functionality, Super I/O, video controller, and remote presence capabilities into a single chip on the compute node system board.

IMM2 provides the following features:

- An embedded web server, which gives you remote control from any supported web browser with the Oracle Java plug-in installed
- A command-line interface (CLI) that the administrator can use from a telnet or SSH session
- Secure Sockets Layer (SSL) and Lightweight Directory Access Protocol (LDAP)
- Multiple alert functions to warn systems administrators of potential problems through email, IPMI Platform Event Traps (PETs), and Simple Network Management Protocol (SNMP)

IMM2 has many enhancements to the previous generation of IMM:

- Faster processor and more memory, delivering a more responsive user interface
- Improved system power-on and boot time
- New web user interface that provides a common design across all Lenovo System x products
- Increased color depth and resolution for remote presence
- Configuration changes require no reboot and become effective immediately
- Hardware management of non-volatile storage
- More information for events detected by the Unified Extensible Firmware Interface (UEFI)
- Simplified update and flashing mechanism
- Security enhancements, including authentication policies

5.1.2 IMM2 Initial setup

The IMM2 provides a virtual presence capability for remote server management through industry-standard interfaces:

- Intelligent Platform Management Interface (IPMI) Version 2.0
- Simple Network Management Protocol (SNMP) Version 3
- Common Information Model (CIM)
- Web browser

The remote presence gives you the ability to perform the following tasks:

- Remotely view video with graphics resolutions up to 1600x1200 at 75 Hz with up to 23 bits per pixel color, regardless of the system state
- Remotely access the server by using the keyboard and mouse from a remote client
- Map the optical drive, diskette drive, ISO, and diskette image files or attached USB flash drives on a remote client for use by the server
- Upload a diskette image to the IMM2 memory and map it to the server as a virtual drive

The following sections provide steps to access the IMM2 and explain how you can configure the IMM2 network interface.

Network interface

You can configure the IMM2 network interface IP address through either the Chassis Management Module (CMM) web interface or CLI.

By default, the IMM2 uses Dynamic Host Configuration Protocol (DHCP) for its management address. If a DHCP host is not available, the IMM2 sets a static IP address, based on which node bay the compute node is installed in. Standard-density nodes' default addresses have a range of 192.168.70.101 to 192.168.70.114.

For instructions on how to assign IP addresses to the IMM, see "Component IP Configuration option" on page 160.

IP addressing for the interface can be set to a static IP, DHCP, or DHCP with static failover. If the CMM and compute nodes are set to use DHCP and the CMM loses connectivity to the network, the CMM attempts to acquire a new IP address when network connectivity is restored. Compute nodes will not attempt to acquire a new IP address and will use their existing IP addresses. Therefore, if you experience network problems on compute nodes after connectivity between the CMM and the network has been restored, you might need to reset the system-management processor on each of the compute nodes in the chassis (including Flex System Manager management software, if it is installed) to trigger acquiring a new DHCP assigned address.

The IMM2 user name and password are provisioned through the CMM. To use tools such as Advanced Settings Utility™ (ASU) for out-of-band communications, you must define local user accounts from the IMM web interface. To define local user accounts from the IMM web interface, click **IMM Management** → **Users**.

The local user accounts are valid only for accessing the IMM web interface on this compute node. However, if you back up the IMM configuration, that backup will not contain the local user accounts that were created from the IMM user interface.

Part of the information that is stored in the IMM2 can be accessed within the UEFI F1 Setup menu by selecting **System Settings** → **Integrated Management Module**. Figure 5-1 on page 74 shows the first panel of the IMM2 configuration panel.

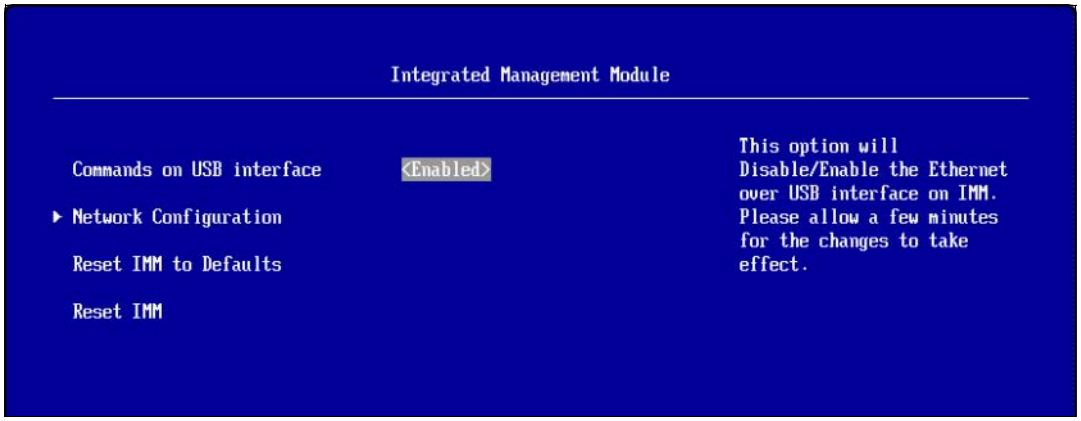

Figure 5-1 Integrated Management Module panel

IMM2 network access

The greatest strength of the IMM2 is the ability to completely monitor and manage the server remotely over the network. How much functionality you have through this remote access depends entirely on your configuration of the IMM2.

IMM2 default configuration

The default network connection for the IMM2 on the X6 Compute Node is through the level 2 switch functionality of the CMM. The IMM management GUI can be launched in-context from the CMM GUI by using your CMM credentials to log in to the IMM GUI or opening a supported web browser to the IP address set for the IMM. An SSH session can be established for CLI-based interaction. The following are the default settings of the IMM2 from the factory:

- Network IP: DHCP. If DHCP fails, use the following static address:
  - IP Address: 192.168.70.1*xx* (where xx is the node's first bay position)
  - Subnet mask: 255.255.255.0
  - Gateway: 0.0.0.0
- Default user ID: `USERID`
- Default password: The current CMM password for the USERID account

5.1.3 Accessing the remote control feature in the IMM2

Follow these steps to use the Remote Control feature on the IMM2 interface for the X6 system:

1. Log in to the IMM2 of the specific system that you want to control.

2. On the System Status page, select **Remote Control,** as shown in Figure 5-2.

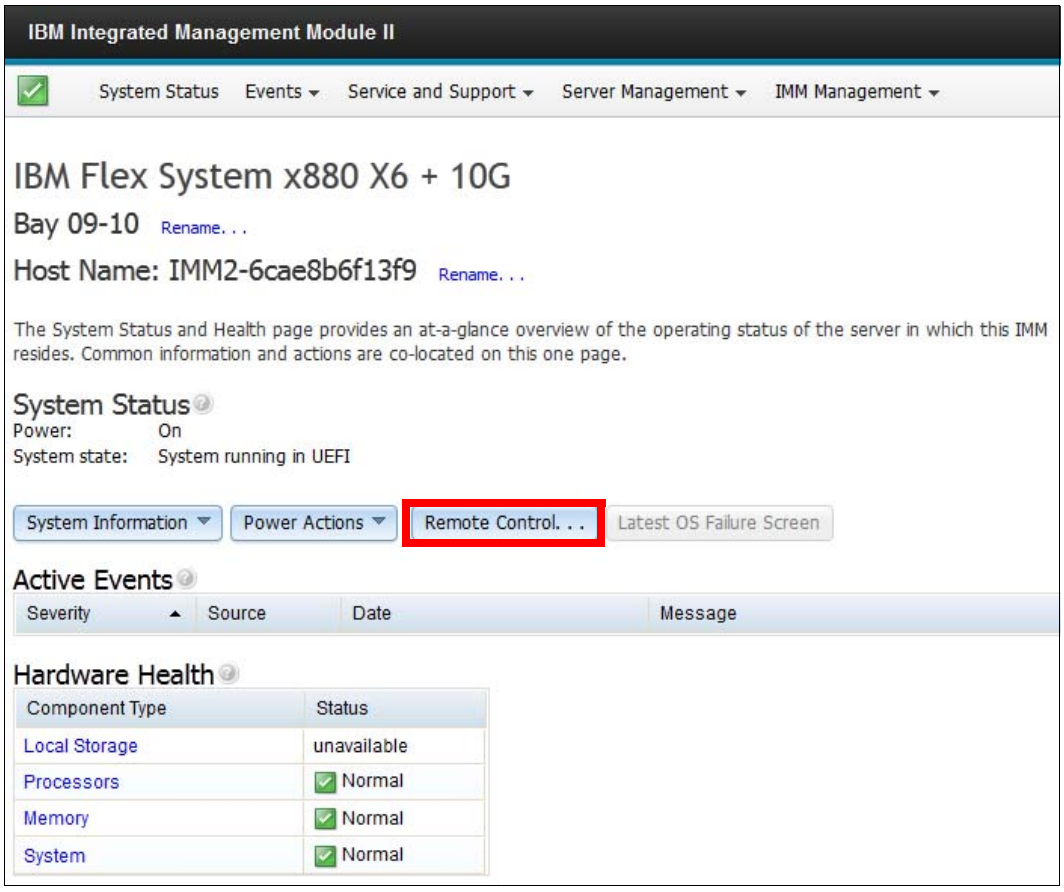

Figure 5-2 Remote control from the IMM2

Figure 5-3 on page 76 shows the Remote Control page in the IMM2 web page.

3. To protect sensitive disk and KVM data during your session, click the **Encrypt disk and KVM data during transmission** check box before starting Remote Control. For complete security, use Remote Control in conjunction with SSL. You can configure SSL at **IMM2 Management** → **Security** from the top menu

4. If you want exclusive remote access during your session, click **Start Remote Control in Single User Mode**. If you want to allow other users remote console (KVM) access during your session, click **Start Remote Control in Multi-user Mode**, as shown in Figure 5-3.

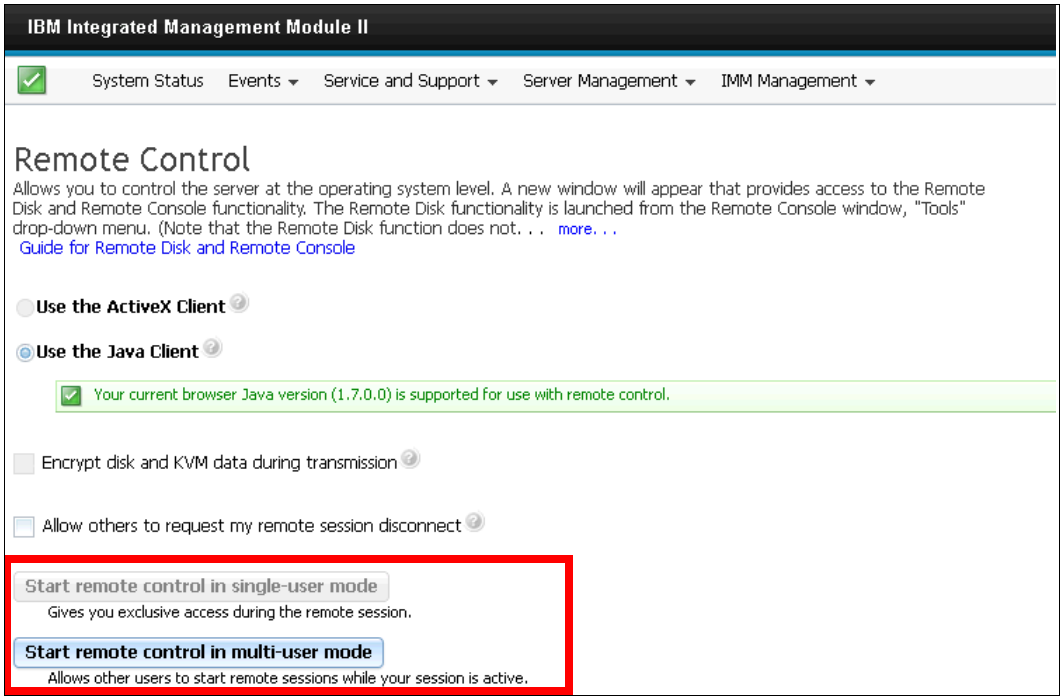

Figure 5-3 IMM2 Remote Control page

5. The various controls that you can use to control the server are described in the *Integrated Management Module II User's Guide*:

   http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086346

Remotely controlling your server will also give you the ability to remote mount an ISO image. See “IMM2 remote media” on page 76 for additional information on how to remote mount an ISO image to your server.

IMM2 remote media

A remote control feature is available through the IMM2 web interface. You must log in to the IMM2 with a user ID that has Supervisor access. You can also assign to the server an optical drive, USB flash drive, or disk image that is on your computer.

For more information, see the *Integrated Management Module II User's Guide* that we cited previously.

Follow these steps to mount a drive through IMM2:

1. Connect to the IMM2 with your web browser.

2. Click the **Remote Control** button on the main page of IMM2, as shown in Figure 5-4:

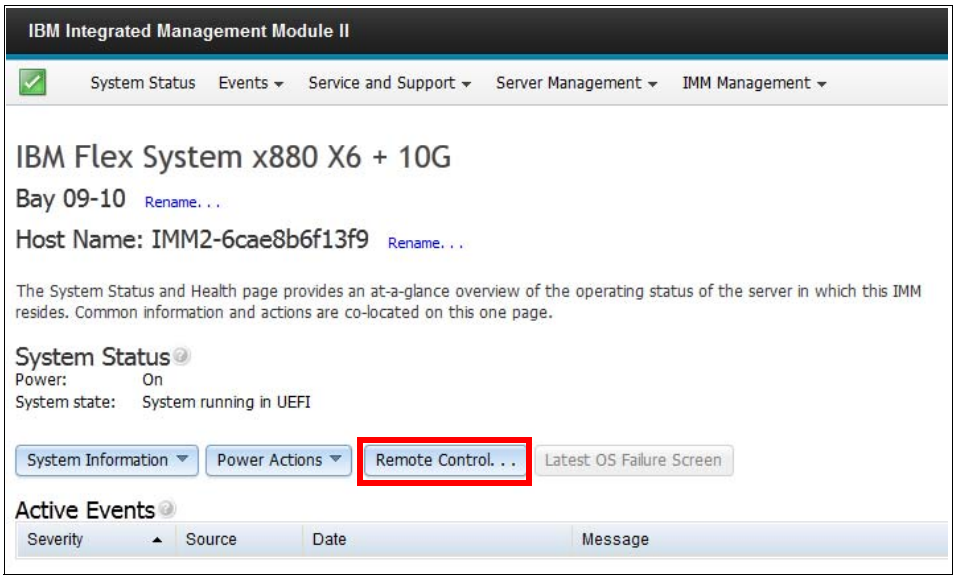

Figure 5-4 Remote Control button on the IMM2 main page

3. If you want to allow other users remote control access during your session, click the **Start Remote Control in Multi-user Mode** button. Otherwise, click **Start Remote Control in Single User Mode**.
4. A Java application window should open, as shown in Figure 5-5:

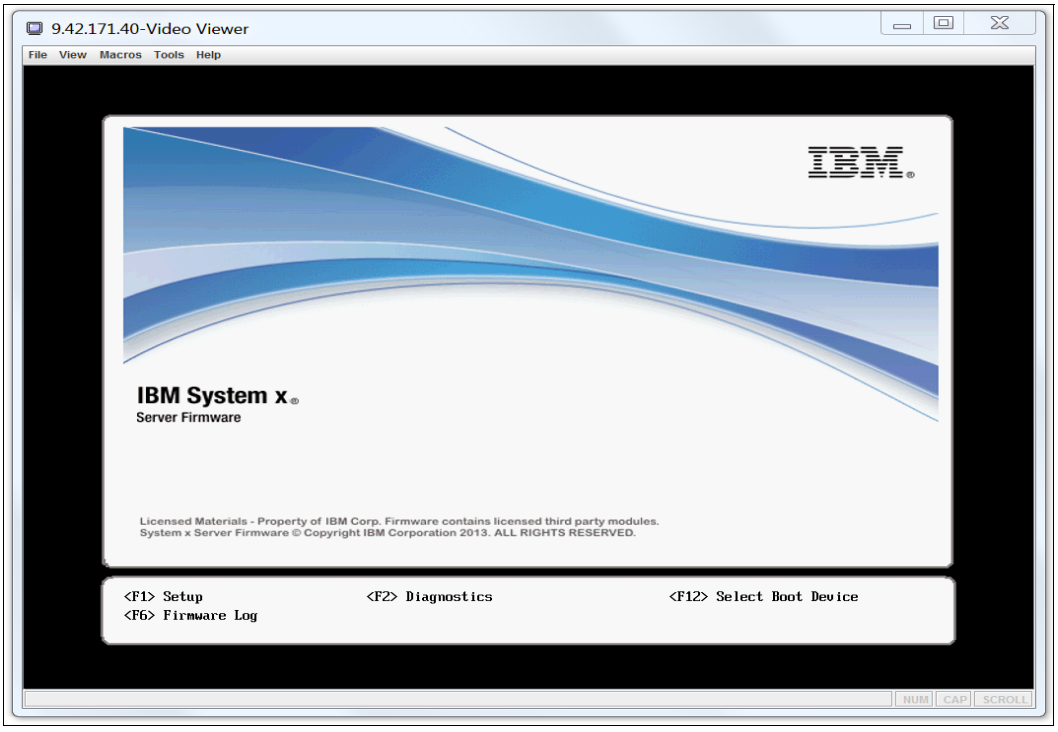

Figure 5-5 Video Viewer window

5. To mount an existing image to the remote server as virtual media, enable the Virtual Media function by clicking **Activate** on the Tools drop-down menu that is shown in Figure 5-6.

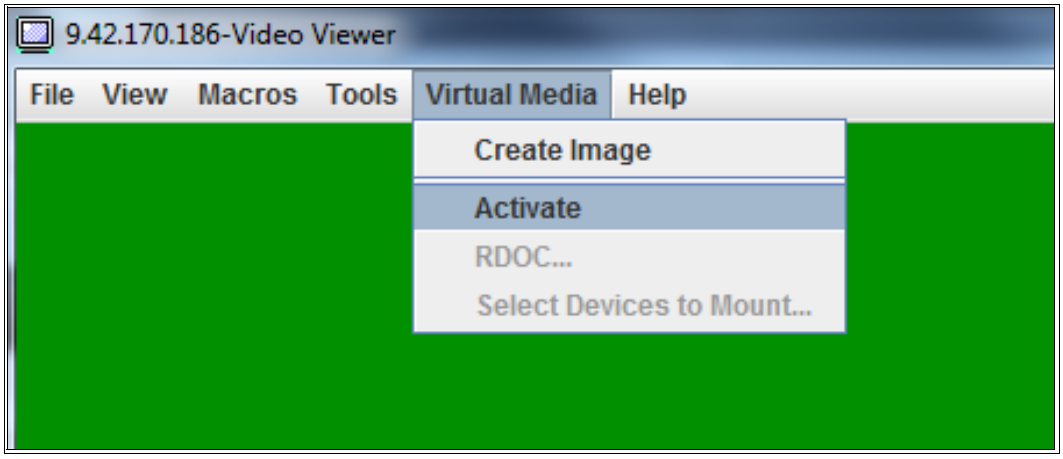

Figure 5-6 Launching virtual media

6. Now, with virtual media active, from the Virtual Media drop-down menu, click **Select Devices to Mount** to bring up the Select Devices to Mount panel shown in Figure 5-7.

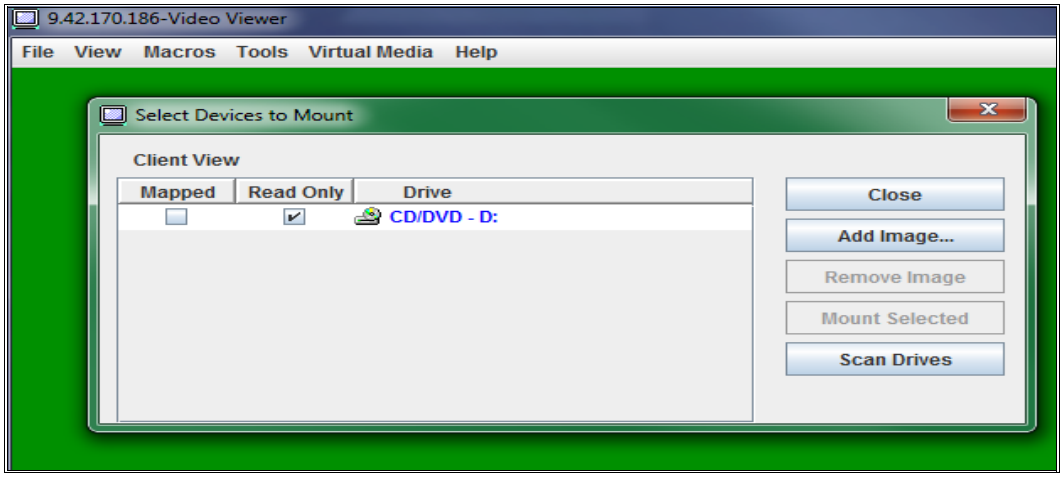

Figure 5-7 The Select Devices to Mount panel

7. After adding the image that you want, select the check boxes next to the drives that you want to map, and click **Mount Selected**, as shown in Figure 5-8:

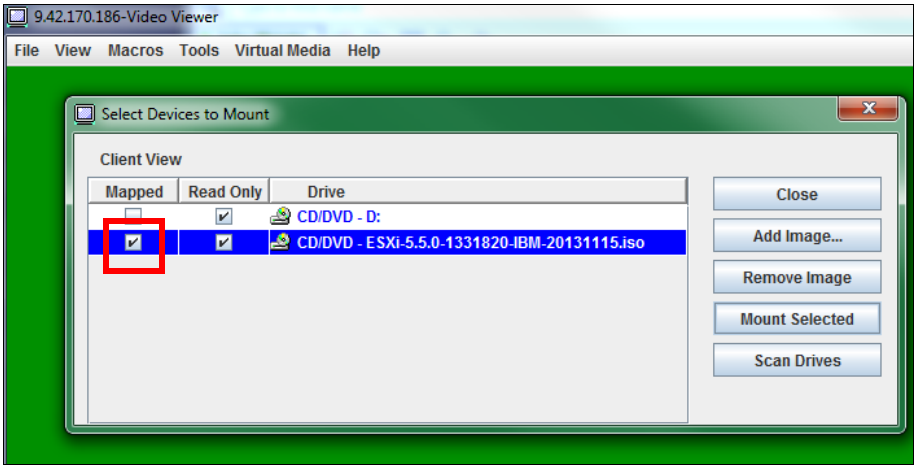

Figure 5-8 Selected image mounting

8. The Select Devices to Mount panel will close, and the mounted devices will show under the Virtual Media drop-down menu shown in Figure 5-9.

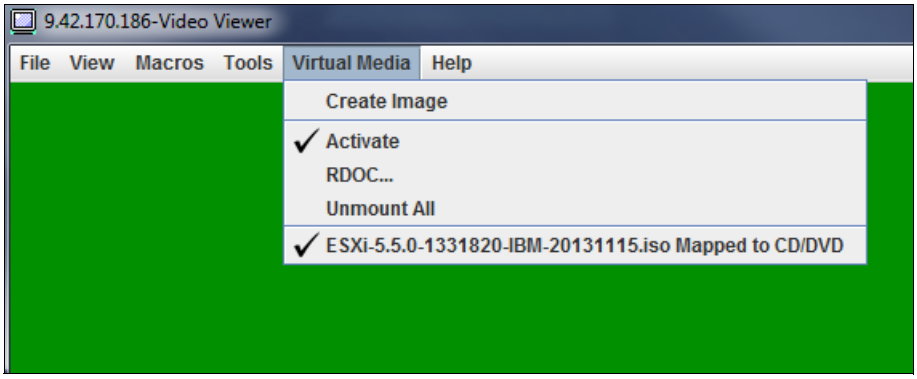

Figure 5-9 Mounted devices shown under the Virtual Media tab

9. IMM Virtual Media-mounted files will show up under Boot Devices Manager within UEFI as bootable devices, as shown in Figure 5-10.

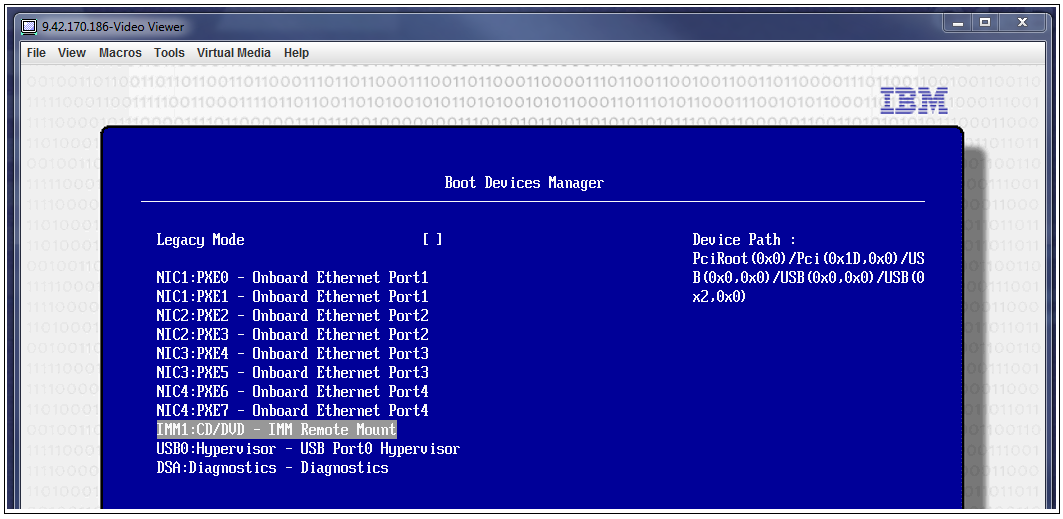

Figure 5-10 Mounted Virtual Media displayed in the UEFI Boot Devices Manager

5.1.4 Initial partition setup

When the X6 complex has been discovered by the CMM and the IMM has been configured, an initial partition in the complex should be set up. This is recommended so that no nodes that are in use in a complex but in an unassigned state are accidentally assigned to a partition with another node within the complex. Ideally, all nodes should be assigned to a partition (even if this is a single whole system partition within the complex) to avoid assigning nodes that appear unassigned but are actually use.

For information on partitioning, see 5.5, “Scalability and partitioning” on page 100.

5.2 Initial setup of UEFI and hardware components

This section provides an overview of the system components and options available for configuration and system optimization when setting up an X6 Compute Node for the first time. It describes the following actions:

- How to change the settings in the UEFI and the different ways to apply these changes.
- How the System Configuration and Boot Management menu structures can be used to configure hardware, such as the RAID controllers and advanced memory modes.
- Where to find and how to select the configurations when you choose to implement a system to improve performance or power use.

Note: Some settings within UEFI, such as CPU throttling or power-saving settings, overrule CMM energy management settings, as required by the chassis. For example, if power supply failures occur, CPUs may be throttled to permit all servers and I/O modules in the chassis to remain online.

This section covers the following topics:

5.2.1 System settings in the Unified Extensible Firmware Interface

The UEFI is the interface between the operating system (OS) and platform firmware. UEFI provides a modern, well-defined environment for booting an OS and running pre-boot applications.

UEFI is effectively the replacement for BIOS. BIOS has been the standard input/output system for many years but was not designed to handle the amount of hardware that can be added to a modern server today. New Lenovo System x models implement UEFI to take advantage of advanced features. See the UEFI website for more information:

http://www.uefi.org/home/

The UEFI page is accessed by pressing **F1** when the System x Server Firmware splash window is displayed, as shown in Figure 5-11.

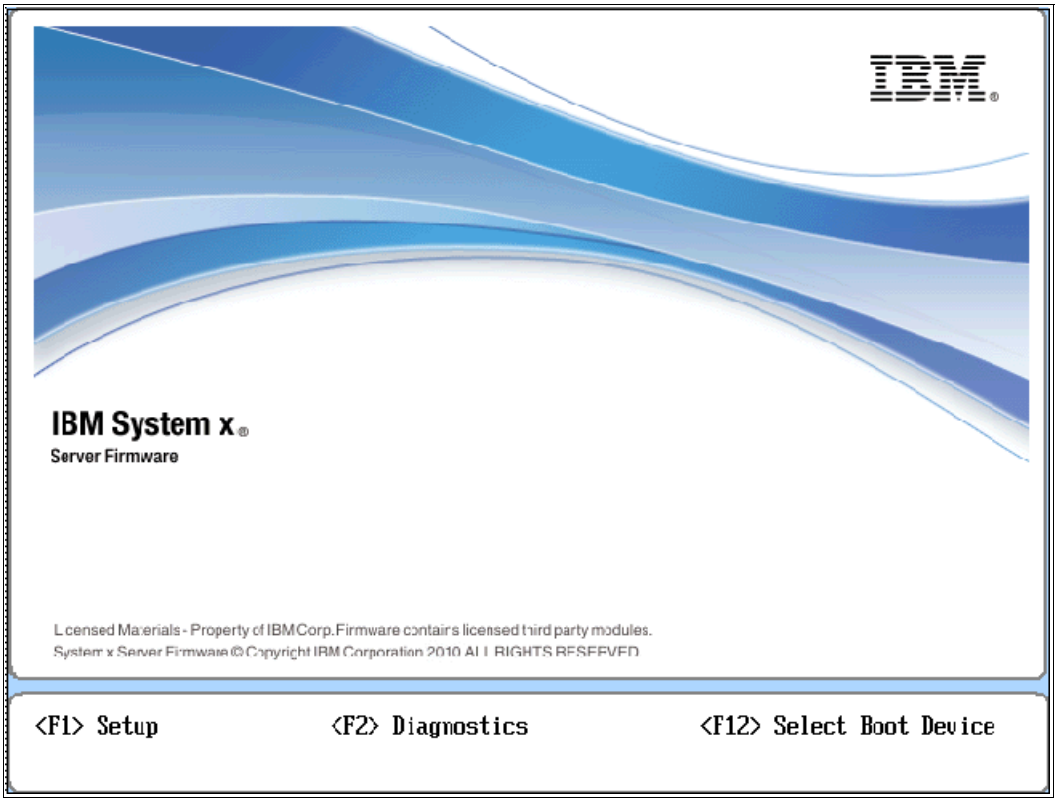

Figure 5-11 Lenovo System x Server Firmware splash screen

You can change UEFI settings to meet your system requirements. In this section, we provide an overview of the UEFI settings for tuning your system for performance.

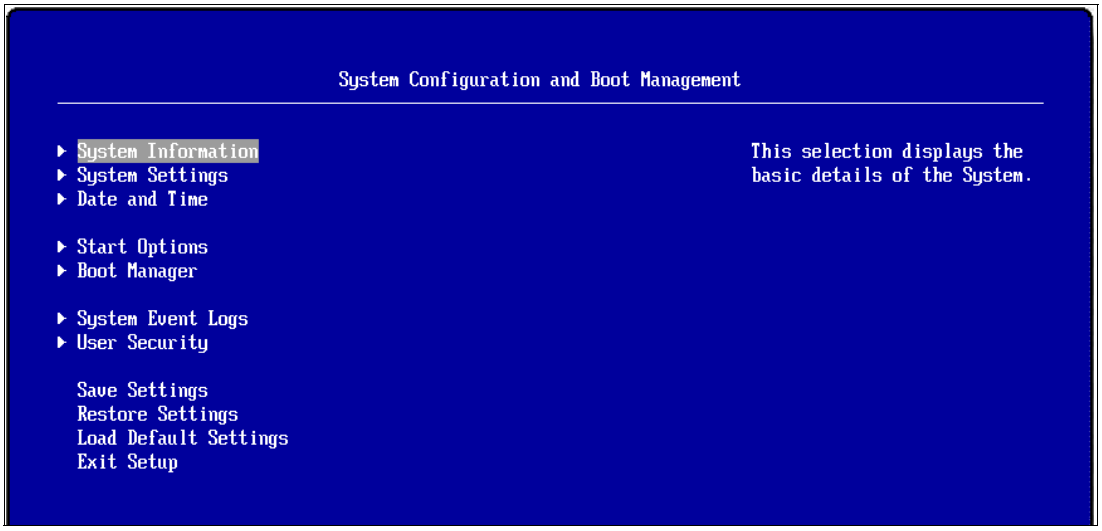

Figure 5-12 UEFI initial panel

Choose **System Settings** to access the settings options that we will describe here, as shown in Figure 5-13

You can use the Advanced Settings Utility (ASU) tool to change UEFI values. ASU exposes more than the settings accessed using the F1Setup panel. For more information, see 7.7, “Advanced Settings Utility (ASU)” on page 172.

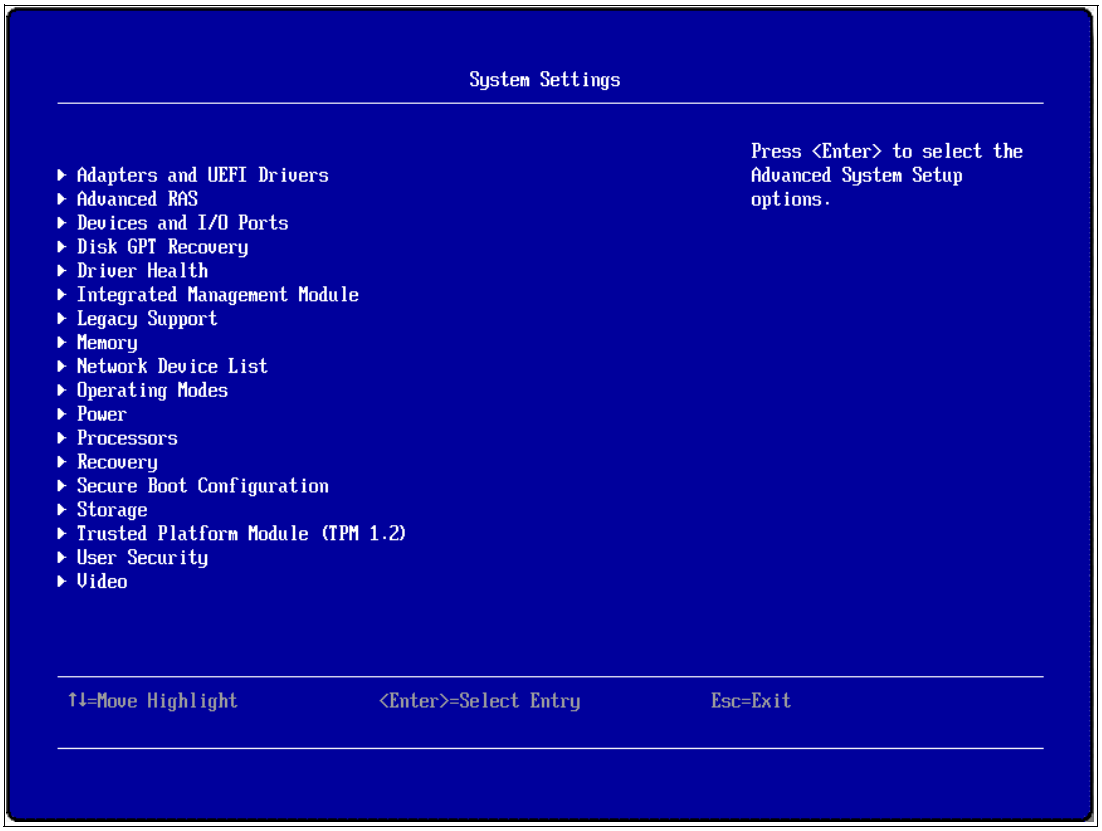

Figure 5-13 UEFI System Settings panel

Table 5-1 provides an overview of the most common UEFI performance optimization settings and their default settings.

Table 5-1 UEFI default settings

UEFI value	Default setting
Processor settings	
Turbo Mode (Intel Turbo Boost)	Enabled
Processor Performance States	Enabled
CPU C-States	Enabled
C1 Enhanced Mode	Enabled
Hyper-Threading	Enabled
Execute Disable Bit	Enabled
Intel Virtualization Technology	Enabled
QPI Link Frequency	Max Performance
Memory settings	
Memory Speed	Max Performance
Patrol Scrub	Enabled
Memory Data Scrambling	Enabled
Advanced RAS	
Machine Check Recovery	Enabled
PCI Live Error Recovery	Enabled

Tuning the X6 Compute Node for performance is a complex topic, because there are many factors, such as which application you have installed or what type of workload your application generates. For example, a database server generates a different load on the hardware than a file and printer server.

In this section, we provide general settings for the X6 Compute Node as a starting point for performance tuning.

Table 5-2 Overview of UEFI settings for specific performance

Setting	Maximum performance	Virtualization[1]	Low latency	Performance per watt	HPC
TurboMode[2]	Enabled	Enabled	Disabled	Disable	Disabled
C states	Disabled	Enabled	Disabled	Enabled	Enabled
C1 Enhanced Mode	Disabled	Enabled	Disabled	Enabled	Enabled
Prefetcher	Enabled	Enabled	Enabled	Enabled	Enabled
Hyper-Threading	Enabled	Enabled	Disabled	Enabled	Enabled
Execute Disable	Disabled	Enabled	Disabled	Enabled	Disabled
Virtualization extensions	Disabled	Enabled	Disabled	Enabled	Disabled

Setting	Maximum performance	Virtualization[1]	Low latency	Performance per watt	HPC
QPI link speed	Maximum performance	Maximum performance	Maximum performance	Power efficiency	Maximum performance
Memory speed	Maximum performance	Maximum performance	Maximum performance	Maximum performance	Maximum performance
Patrol scrub	Disabled	Disabled	Disabled	Disabled	Disabled

1. These virtualization settings are recommended for a stand-alone host only. For multiple virtualized hosts in clustered workloads, use the *Maximum performance* settings instead.
2. Depending on the processor workload, enabling TurboMode might also increase power consumption. The actual processing performance boost that you get from TurboMode depends on the temperature and humidity in environment that the server is in, because the processor will boost performance only up to the environmental limits set by the processor design.

5.2.2 UEFI common settings

The UEFI default settings are configured to provide optimal performance with reasonable power consumption. Other operating modes, as described in 2.6, “Power policies” on page 18, are also available to meet various power and performance requirements. However, individual system settings enable you to fine-tune the characteristics of the Lenovo X6 Compute Nodes.

This subsection describes these individual system settings:

- “System power settings”
- “Processor settings”
- “Memory settings” on page 86

System power settings

Power settings that are listed in the UEFI interface are overridden by what is configured within the CMM. For CMM power control settings, see 2.6, “Power policies” on page 18.

Processor settings

Processor settings control the various performance and power features that are available on the installed Xeon processor.

Figure 5-14 shows the UEFI Processors system settings window with the default values.

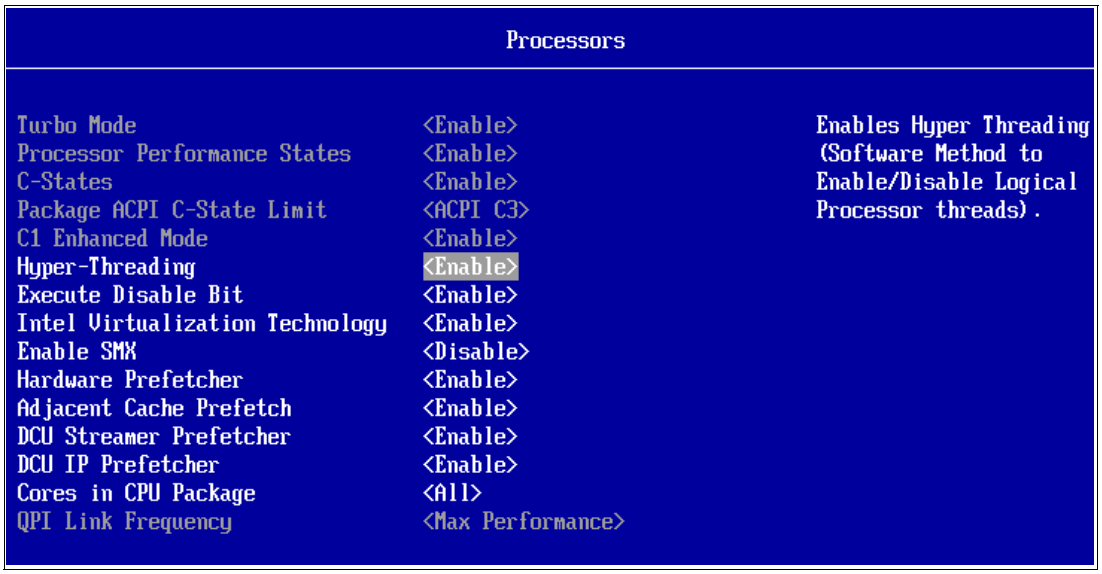

Figure 5-14 UEFI Processor settings panel

The following processor feature options are available:

- Turbo Mode (Default: Enabled)

  This mode enables the processor to increase its clock speed dynamically if the CPU does not exceed the Thermal Design Power (TDP) for which it was designed.

- Processor Performance States (Default: Enabled)

  This option enables Intel Enhanced SpeedStep Technology that controls dynamic processor frequency and voltage changes, depending upon operation.

- CPU C-States (Default: Enabled)

  This option enables dynamic processor frequency and voltage changes in the idle state, which potentially conserves power.

- Package ACPI C-State Limit (Default: ACPI C3)

  The Advanced Configuration and Power Interface (ACPI) is configured as ACPI C3 by default.

- C1 Enhanced Mode (Default: Enabled)

  This option enables processor cores to enter an enhanced halt state to lower the voltage requirement, which conserves power.

- Hyper-Threading (Default: Enabled)

  This option enables logical multithreading in the processor so that the operating system can execute two threads simultaneously for each physical core.

- Execute Disabled Bit (Default: Enabled)

  This option enables the processor to disable the execution of certain memory areas to prevent buffer overflow attacks.

- Intel Virtualization Technology (Default: Enabled)

  This option enables the processor hardware acceleration feature for virtualization.

- Enable SMX (Default: Disabled)

  This option enables Safer Mode Extensions.

- Hardware Prefetcher (Default: Enabled)

  This option enables hardware prefetching. Lightly threaded applications and some benchmarks can benefit from having the Hardware Prefetcher enabled.

- Adjacent Cache Prefetcher (Default: Enabled)

  This option enables adjacent catch prefetching. Lightly threaded applications and some benchmarks can benefit from having the Adjacent Cache Prefetcher enabled.

- DCU Streamer Prefetcher (Default: Enabled)

  This option enables DCU streamer prefetching. Lightly threaded applications and some benchmarks can benefit from having the DCU Streamer Prefetcher enabled.

- DCU IP Prefetcher (Default: Enabled)

  This enables the DCU IP Prefetcher. It is best to leave this enabled for most environments. Some environments might benefit from having it disabled (for example, Java).

- Cores in CPU Package (Default: All)

  This option sets the number of processor cores to be activated within each CPU package. You might want to change your CPU cores to lower your power consumption or to meet software licensing requirements.

- QPI Link Frequency (Default: Max Performance)

  This option sets the operating frequency of the processor's QPI links:

  - *Minimal Power* provides less performance for better power savings.
  - *Power Efficiency* provides the best performance per watt ratio
  - *Max Performance* provides the best system performance

Memory settings

The Memory settings panel provides the available memory operation options, as shown in Figure 5-15.

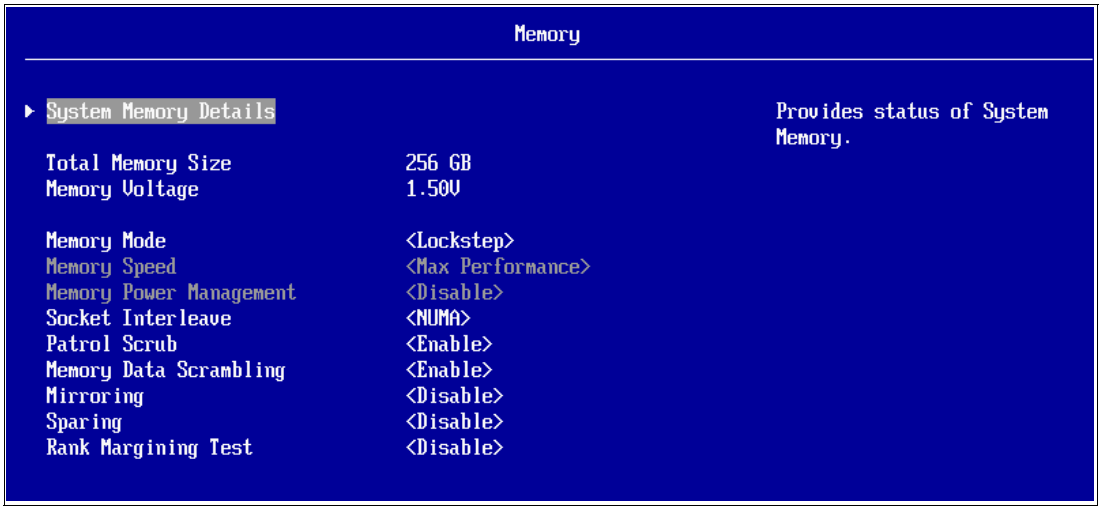

Figure 5-15 UEFI Memory settings panel

The Memory settings panel provides the following options:

- Memory Mode (Default: Lockstep)

  This option enables Lockstep (RAS) or Independent memory mode. In Lockstep mode, two channels operate synchronously, allowing advanced memory protection, with SMI link operations at the DDR3 speed. With Independent mode, all channels operate

independently, with SMI links operating at twice the DDR3 speed. For more information on Memory Modes, see 5.3, “Configuring ServeRAID M1200 using UEFI” on page 87.

- Memory Speed (Default: Max Performance)

  This option sets the operating frequency of the installed DIMMs:

  - *Minimal Power* provides less performance for better power savings. The memory operates at the lowest supported frequency, which is 1333 MHz in the x6 Compute Node.
  - *Power Efficiency* provides the best performance per watt ratio. The memory operates one step under the rated frequency.
  - *Max Performance* provides the best system performance. The memory operates at the rated frequency, that is, 1600 MHz for DIMMs rated at 1600MHz or higher.

- Memory Power Management (Default: Disabled)

  This option enables or disables Memory Power Management.

- Socket Interleave (Default: NUMA)

  This option sets Socket Interleave to NUMA (Non-Unified Memory Architecture) or Non-NUMA. In NUMA mode, memory is *not* interleaved across processors. Non-NUMA mode memory is interleaved across processors.

- Patrol Scrub (Default: Enabled)

  This option enables scheduled background memory scrubbing before any error is reported, as opposed to default demand scrubbing on an error event. This option provides better memory subsystem resiliency at the expense of a small performance loss.

- Memory Data Scrambling (Default: Enabled)

  This option enables a memory data scrambling feature to further minimize bit-data errors.

- Mirroring (Default: Disabled)

  This option enables or disables memory mirroring as described in 3.6.2, “Memory mirroring” on page 38.

  Mirroring: Memory Mirroring mode cannot be used in conjunction with Memory Sparing mode (rank sparing).

- Sparing (Default: Disabled)

  This option enables or disables memory rank sparing as described in 3.6.3, “Rank sparing” on page 39.

5.3 Configuring ServeRAID M1200 using UEFI

For the ServeRAID M1200 SAS/SATA controller, you can configure and monitor arrays and attached devices from the UEFI interface.

The UEFI provides access to the LSI MegaRAID Configuration Utility that is used by the ServeRAID M1210e SAS/SATA Controller. Access the menu in UEFI by selecting **System Settings** → **Storage** → **LSI MegaRAID <ServeRAID M1210e> Configuration Utility**, as shown in Figure 5-16.

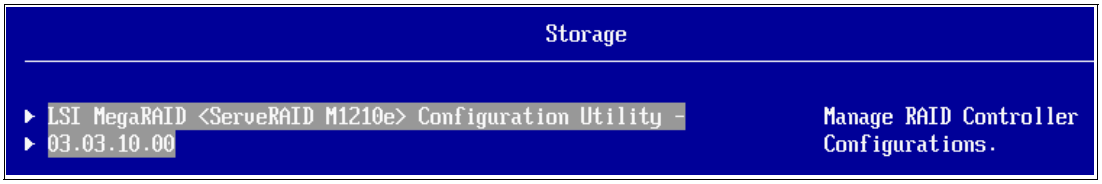

Figure 5-16 LSI MegaRAID Utility panel for ServeRAID M1210e

As Figure 5-17 shows, the LSI MegaRAID Configuration Utility panel Main Menu offers several options for managing your internal disks.

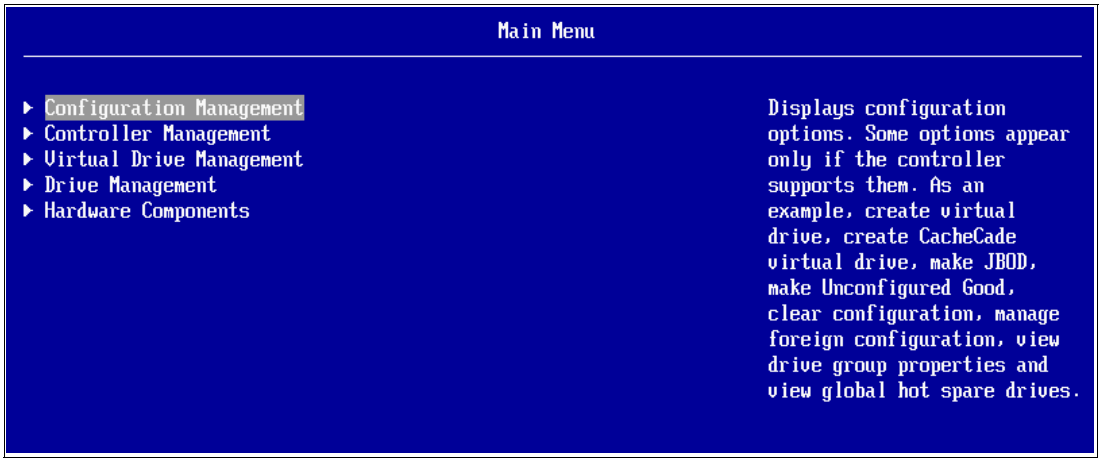

Figure 5-17 Configuration Utility, Main Menu panel

This menu includes the following settings:

- Configuration Management

  Displays configuration options. Some options, such as those that follow, appear only if the controller supports them:

  - Clear Configuration
  - Create CacheCade Virtual Drive
  - Create Virtual Drive
  - Make JBOD
  - Make Unconfigured Good
  - Manage Foreign Configuration
  - View Drive Group Properties
  - View Global Hot Spare Drives

- Controller Management

  Displays the controller status and basic properties of the controller, such as product name, serial number, PCI ID, firmware version, NVDATA version, reset to factory defaults, and Clear Configuration. You can also use the advanced link to view additional properties and perform additional tasks, such as changing the security key or saving the TTY log.

- Virtual Drive Management

  Manages the virtual drive properties and enables you to view the basic virtual drive properties and perform operations, such as background initialization and consistency check. You can also view additional properties by clicking the Advanced link.

- Drive Management

  Displays the basic drive properties and performs operations such as assign or unassign a hot spare drive, locate drives, place drive offline or online, and rebuild a drive. You can also view additional properties by clicking the Advanced link.

- Hardware Components

  Displays the battery status and the status of temperature sensors, fans, and power supplies. You can also view additional properties and perform additional operations by clicking the Advanced link. Some options will only appear if the controller supports them.

The UEFI supports configuration of your RAID array with a supported controller. The following example guides you through the creation of a RAID-1 array configured in the UEFI:

1. Enter UEFI setup by selecting **F1** when prompted at boot time. Access the LSI Configuration Utility from the menu by selecting **System Settings** → **Storage** → **LSI MegaRAID <ServeRAID M1201e> Configuration Utility**, as shown in Figure 5-18.

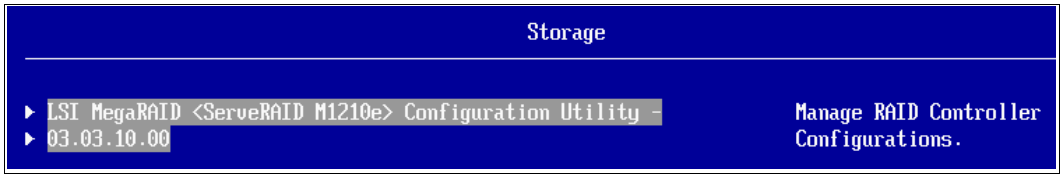

Figure 5-18 Launching the MegaRAID Configuration Utility for the ServeRAID M1210e

2. Figure 5-19 on page 90 shows the Main Menu for the Configuration Utility. Select **Controller Management**.

Figure 5-19 LSI MegaRAID Configuration Utility, Main Menu panel

3. Access the **Controller Management** panel and then select Advanced, as shown in Figure 5-20.

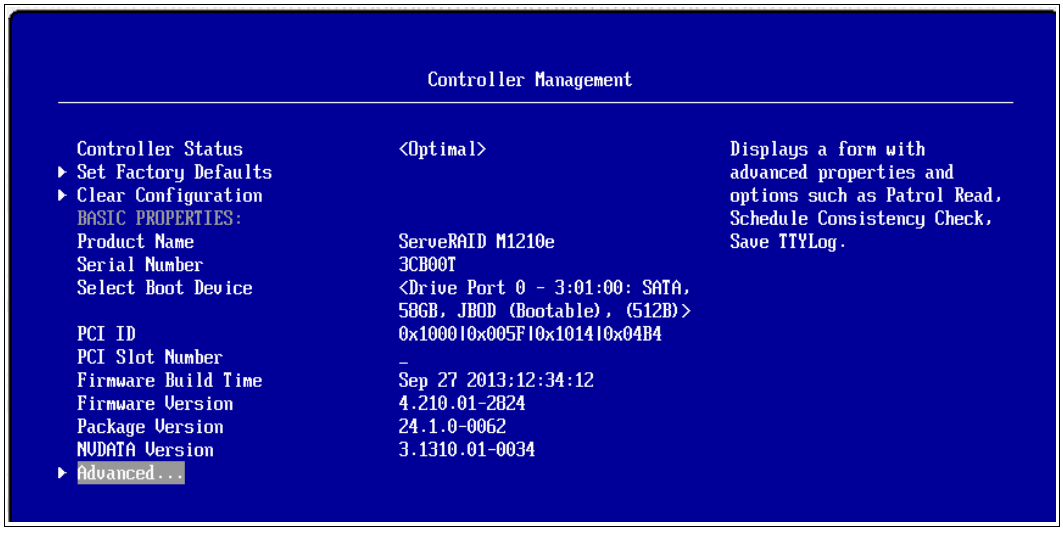

Figure 5-20 Controller Management panel, Advanced option selected

4. From the Advanced Controller Management menu, select **Manage MegaRAID Advanced Software Options**, as shown in Figure 5-21 on page 91.

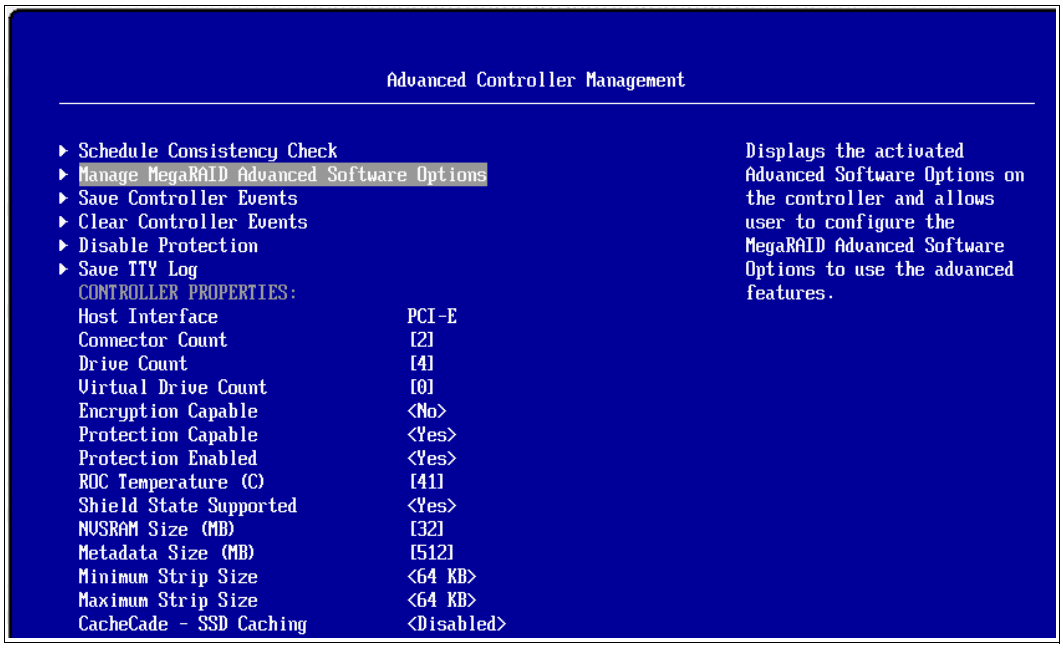

Figure 5-21 Advanced Controller Management panel

5. Ensure the MegaRAID Advanced Software Options is activated as Figure 5-22 shows.

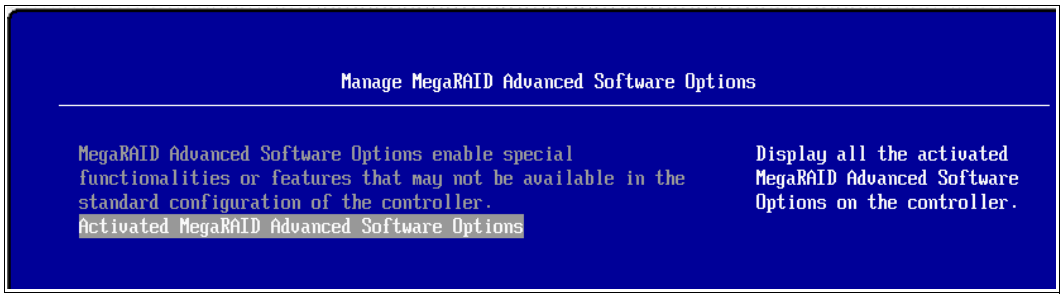

Figure 5-22 Manage MegaRAID Advanced Software Options

6. Exit this menu by pressing the Esc key and return to the **Main Menu** shown in Figure 5-19 on page 90.
7. From the Main Menu, select **Drive Management**. Ensure that your disks are in the *Unconfigured Good* state by selecting each drive individually so you can check the status for each one. For example, the drive shown in Figure 5-23 has the Unconfigured Good status.
8. From each drive window (Figure 5-23 on page 92), highlight **Select operation** and press Enter. Then, select **Initialize Drive** and press Enter to initialize the drive.

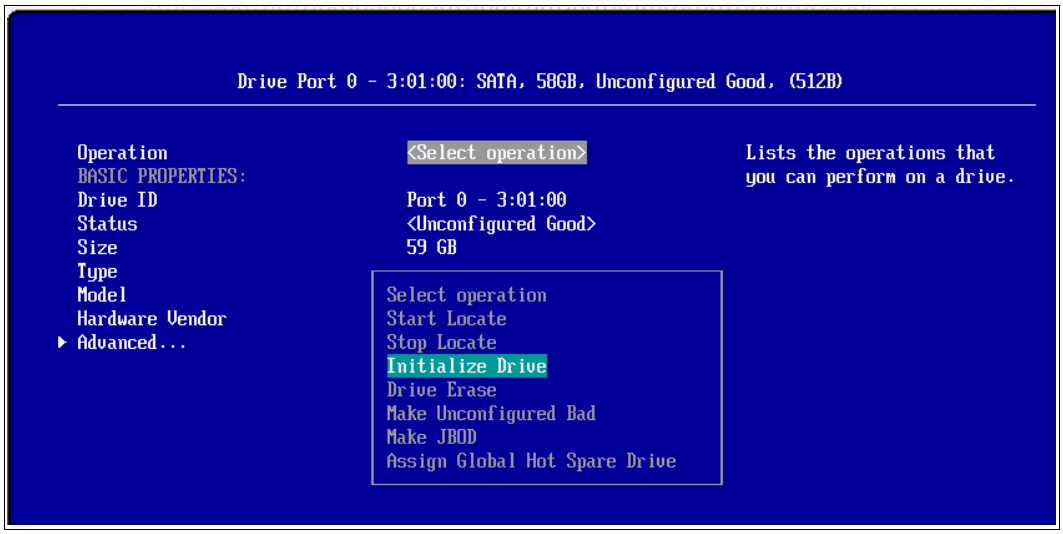

Figure 5-23 Drive operation selection, Initialize Drive selected

9. Return again to the **Main Menu** shown in Figure 5-19 on page 90. Access the **Configuration Management** panel, and select **Create Virtual Drive** as shown in Figure 5-24.

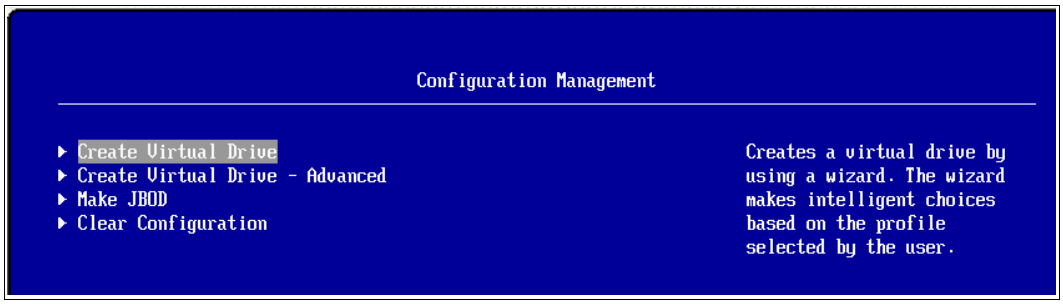

Figure 5-24 Configuration Management panel, Create Virtual Drive selected

10. From the Create Virtual Drive panel, select the RAID level that you want for your array. For this example, we selected RAID 1 as shown in Figure 5-25.

Figure 5-25 Create Virtual Drive panel, Generic RAID 1 selected

11. From the Generic R1 panel, select the **Save Configuration** option to confirm your intentions to create your RAID array (Figure 5-26).

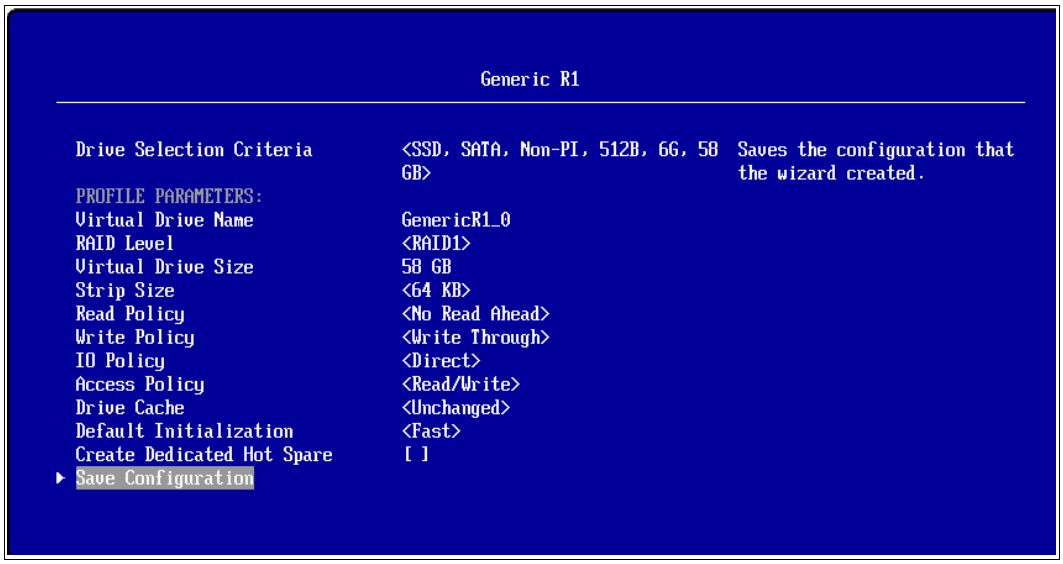

Figure 5-26 Generic RAID 1 array panel

12. You are prompted that creating a RAID array will erase data. Select the **Confirm** option and press the space bar to mark it with an **X**, and then select **Yes**.

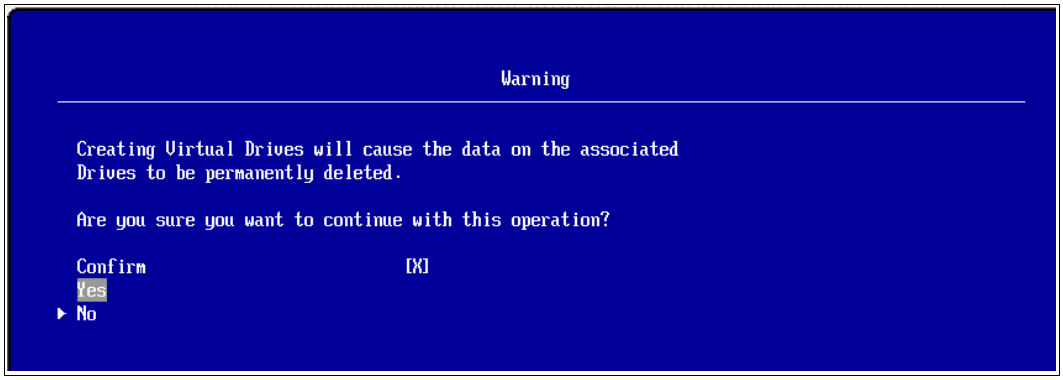

Figure 5-27 Warning and confirmation of RAID array creation

After you confirm your array creation, you will be presented with a Success message as Figure 5-28 shows.

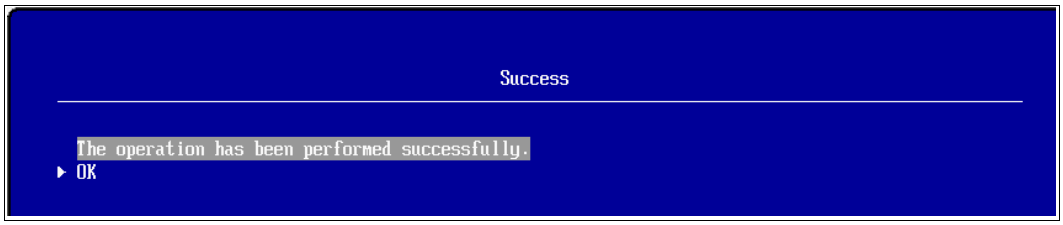

Figure 5-28 Successful RAID array creation confirmation panel

13. Select **OK**, and press Esc to return to the main menu. Your array is now active.

There are more utilities for configuring and managing your disks in the x6 Compute Node. See 6.2.3, “Lenovo ServerGuide” on page 134 and 7.8, “MegaRAID Storage Manager” on page 175.

5.3.1 Drive sharing with vertical RAID

If you have a dual-node X6 complex (two x480 X6 Compute Nodes or two x880 X6 Compute Nodes), the drives in both compute nodes can all be managed by one ServeRAID M1210e embedded controller (the M1210e in the primary or lower compute node). This feature is known as *Drive Sharing* in UEFI.

In such a configuration, you can use all four drives installed in the two nodes to form a single RAID volume. The M1210e in the secondary compute node is not used.

For more details on the M1210e adapter, see 4.9, “Internal disk controller” on page 61.

Note: The vertical RAID feature supports only four drives across two compute nodes. An eight-drive configuration is not supported.

Drive sharing (also known as *vertical RAID*) can, by default, create RAID-0 (JBOD) and RAID-1 (mirrored) arrays. If the M1210e in the primary (lower) compute node is upgraded to support RAID-5 through Lenovo Features on Demand (FoD), you can then form a RAID-5 array using all four drives. An FoD upgrade is not needed in the secondary (upper) compute node in the two-node complex.

To enable the shared drive function of the M1210e, ServeRAID controllers perform the following actions:

1. A dual-node complex or partition with a M1210e in both nodes will show both controllers separately within UEFI, as shown in Figure 5-29, if Drive Sharing is disabled. Browsing both controllers will show them as separate controllers with the physical drives that they are attached to.

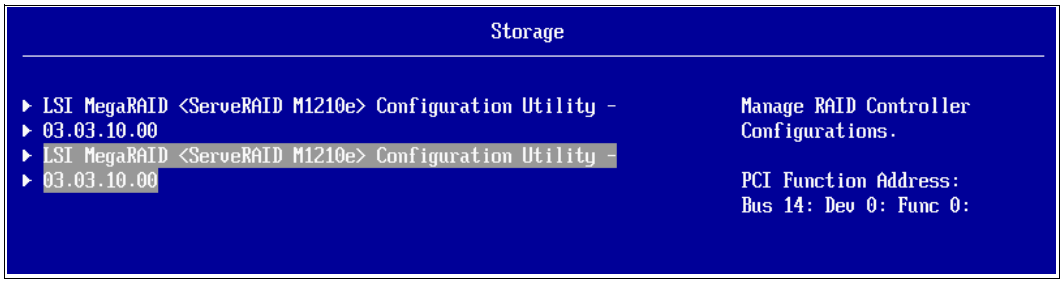

Figure 5-29 Dual-node complex with Drive Sharing disabled

2. Access the menu in UEFI by selecting **System Settings** → **Devices and I/O Ports** as shown in Figure 5-30. Highlight **Disabled** next to Drive Sharing and press Enter to change Disabled to Enabled. Exit UEFI and allow the system to restart.

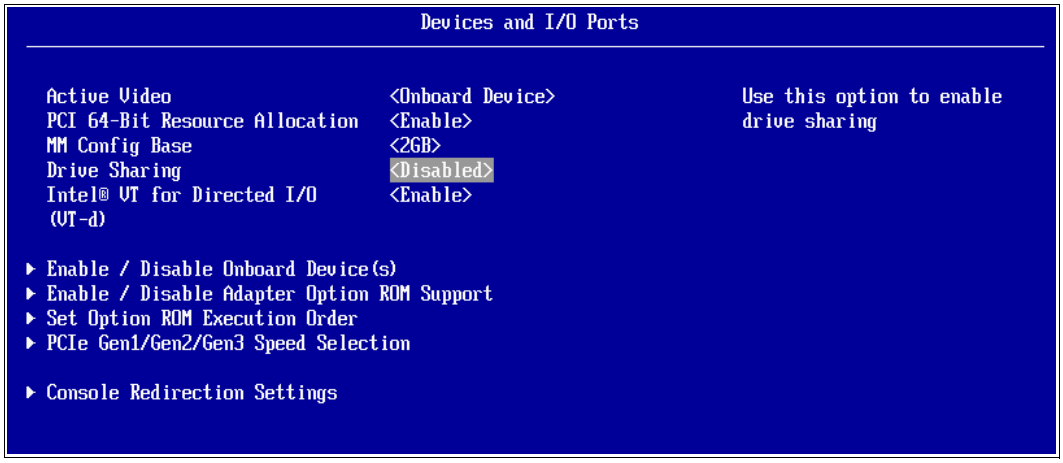

Figure 5-30 Drive Sharing enablement for vertical RAID

3. After the system has restarted, access the UEFI menu again by pressing **F1** at the splash screen and then clicking **System Settings** → **Storage**. As shown in Figure 5-31, only a single M1210e controller is now listed. Browsing the controller will show all physical drives in the complex listed under it. The M1210e controller of the secondary node will not be displayed.

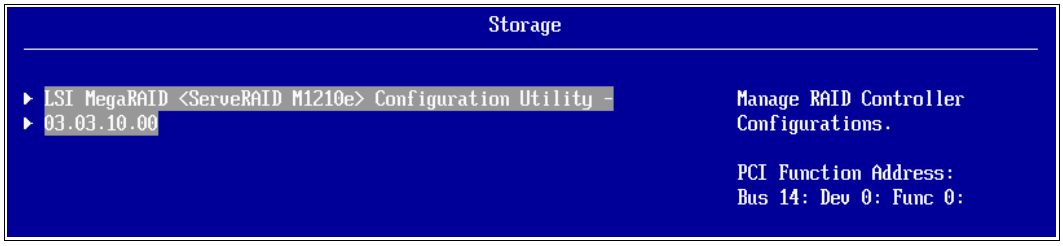

Figure 5-31 Drive Sharing mode enabled displaying primary node raid controller only

4. RAID arrays can now be set up as described in 5.3, “Configuring ServeRAID M1200 using UEFI” on page 87.

Note: RAID arrays can now be configured where data spans two nodes in a partition. Before performing any X6 complex partitioning actions by using the IMM, ensure that data stored on these arrays is backed up, because a loss of data could occur otherwise.

5.4 Enabling Features on Demand

The Lenovo Features on Demand (FoD) function can install and manage optional server and systems management features.

There are two ways of enabling features managed by the FoD process:

- Select them as part of your initial server order and the licenses are installed and enabled as part of the factory build.

- Purchase the FoD options separately and apply them through the IMM2. You can also apply them through other management tools, such as the ToolsCenter™ Activation Management Utility or the Flex System Manager.

This section introduces the process of installing a FoD activation key by using the IMM2:

1. Purchase the FoD key through your normal channels, using the part numbers listed in Table 5-3.
2. An authorization code is sent to you by email or by postal mail within two to five days.
3. Enter the authorization code on the Lenovo FoD web page to generate an activation key file. This process is described in 5.4.1, “How to request an activation key” on page 96.
4. Use the IMM2 web interface or CLI to upload the activation key to the IMM2. This process is described in 5.4.2, “Adding, viewing, and removing an activation key” on page 98.

The FoD function is now enabled on the server. The activation data is saved in nonvolatile random access memory (NVRAM) within the IMM2.

Authorization code: The authorization code that you receive and activate is used on one server. The IMM2 interface also provides a way to remove a feature, if necessary.

These current FoD options for the X6 Compute Node family are listed in Table 5-3:

- Embedded ServeRAID M1210e RAID 5 upgrade

  The embedded RAID controller supports RAID-5 through an FoD upgrade. In a multi-node complex, every active controller needs its own FoD upgrade if you want to use RAID-5 on that controller.

  For example, if you have a two-node complex and you are using Drive Sharing (vertical RAID), then only one M1210e controller is active (the one in the primary or lower node). Therefore, you need only one FoD upgrade.

- CN4054R 10Gb Virtual Fabric Adapter FCoE upgrade

  This adapter supports FCoE or iSCSI through FoD. One FoD upgrade is needed for each adapter installed in a compute node. The FoD upgrade includes two activation codes, one for each ASIC on the adapter.

Table 5-3 Available Feature-on-Demand part numbers

Part number	Feature code	Feature on Demand description	Supported hardware component
00AE930	A5H5	ServeRAID M1200 Zero Cache/RAID 5 Upgrade for Lenovo Systems (FoD)	Embedded ServeRAID M1210e controller
90Y3558	A1R0	CN4054 Virtual Fabric Adapter Upgrade	CN4054R 10Gb Virtual Fabric Adapter, 00Y3306

5.4.1 How to request an activation key

When you purchase an FoD option, you receive this information by email or postal mail:

- Order number
- Feature name and part number
- Quantity ordered
- Authorization Code

Before an activation key can be applied by using the IMM2, it must first be generated by using the Activation Code provided by mail or email at the time of purchase.

You can request an activation key from the Features on Demand web page:

http://www.ibm.com/systems/x/fod

From that page, log in with your user name and password. If you do not already have a Lenovo account, follow the links to create one.

1. Click **Request activation key**. Figure 5-32 shows this step (see selection at left).

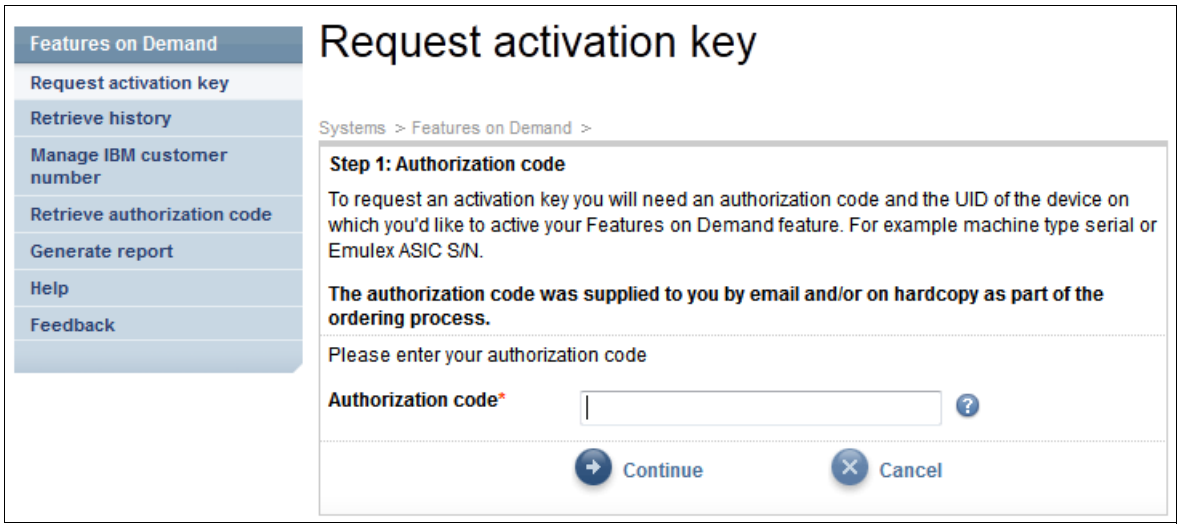

Figure 5-32 Entering the authorization code

2. Enter the authorization code provided and click **Continue** to go to the “Machine details” page shown in Figure 5-33. Provide the machine type and serial number of the server in which you want to activate the feature and click **Generate key**.

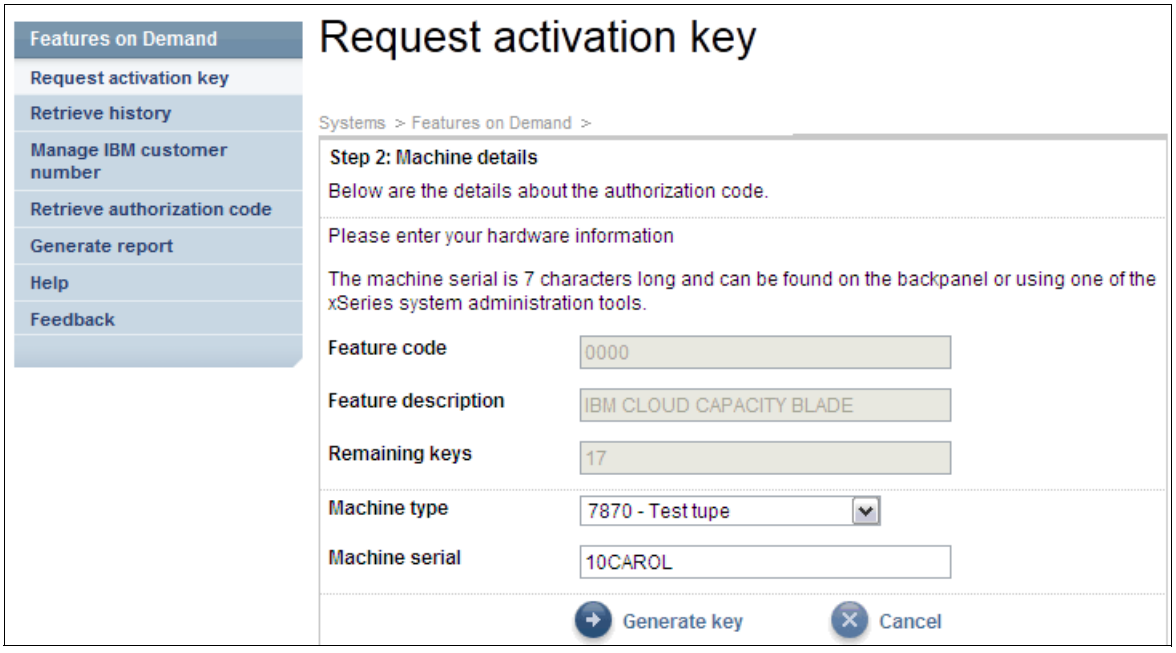

Figure 5-33 Enter the machine type and serial number

3. Verify that the information is correct, and then click **Confirm**.

You can then download an activation key file that is unique to the intended server and can be applied through the IMM2.

You can also use the FoD website to search for previous activation keys that you generated. Figure 5-34 shows how to search for the activation key history.

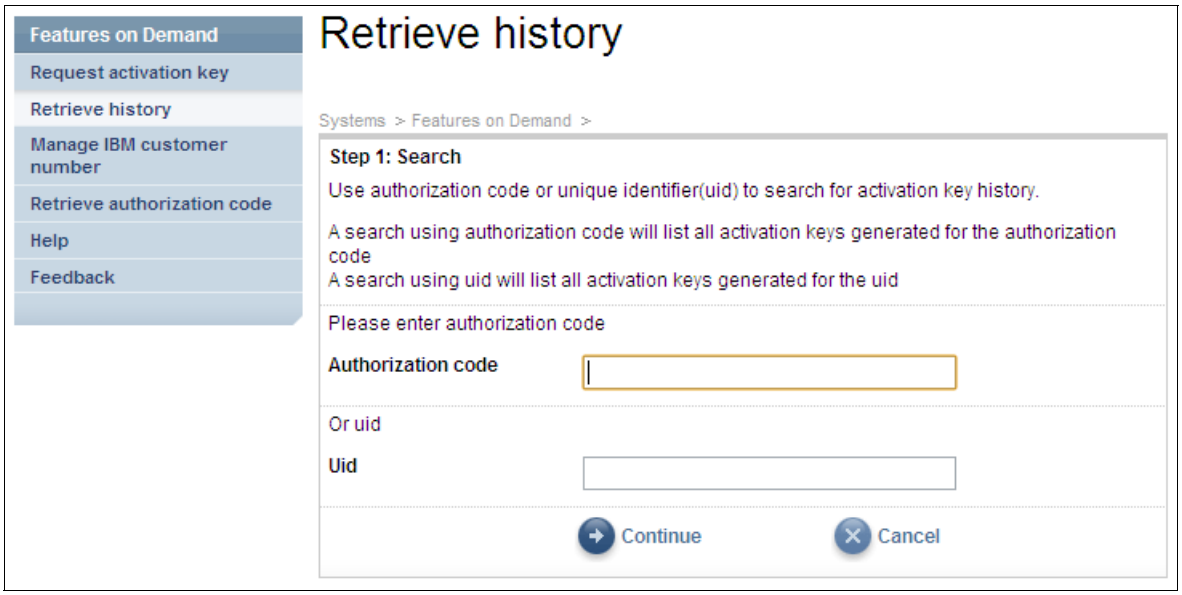

Figure 5-34 Searching for activation keys from the "Retrieve history" page

For more information about FoD, see the *Features on Demand User's Guide avilable from:*

`http://ibm.com/support/entry/portal/docdisplay?lndocid=migr-5089568`

5.4.2 Adding, viewing, and removing an activation key

A FoD activation key can be added and removed from a server through the IMM2 web or CLI interfaces. An activation key can also be deployed by using management tools, such as the ToolsCenter Activation Management Utility or IBM Systems Director Configuration Manager. The IMM2 web interface is described in this section.

Adding an activation key

Before beginning, ensure that you followed the steps described in 5.4.1, "How to request an activation key" on page 96. Have the activation key available from the workstation in which you are running the IMM2 web interface.

Follow these steps to add an activation key:

1. Log in to the IMM2 through a web browser.

2. Select **Activation Key Management** from the IMM Management menu. Figure 5-35 shows this menu.

Figure 5-35 Activation Key Management

3. Under Activation Key Management, click **Add**.
4. In the Add Activation Key window, click **Select File**. Browse to the saved location for the activation key and click **OK**.
5. A pop-up window (Figure 5-36) confirms that the activation key was added successfully.

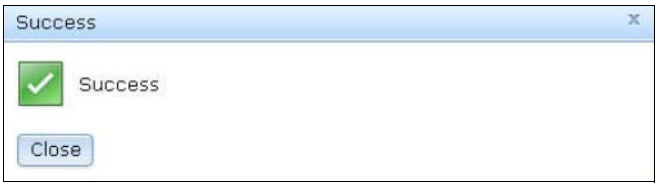

Figure 5-36 FoD successfully enabled

Failure to add the key: An activation key fails to add if it is not intended for the specified machine type or is an invalid code.

Viewing installed keys

Log in to the IMM2 through the web interface by using the steps in “Adding an activation key” on page 98.

Then, select **Activation Key Management** in the IMM Management menu. In this menu, you see a list of current active FoD keys that have been added to the server.

Removing an activation key

The FoD activation keys enabled on the server can also be removed by using the IMM2 web interface. After the key is removed, it no longer shows in the Activation Key Management page of the IMM2. The following steps remove the key from the server:

1. Log in to the IMM2 through a web browser.
2. Select **Activation Key Management** from the IMM Management menu. Figure 5-35 on page 99 shows this menu.
3. Click the feature key to remove. When it is selected, click **Delete**.
4. In the Confirm Activation Key Deletion window, click **OK** to confirm the deletion.

5.5 Scalability and partitioning

The X6 Compute Nodes have a flexible modular design. This enables you to increase the server's compute power and I/O capabilities by adding additional compute nodes to your X6 Compute Node complex.

5.5.1 Overview

This section includes the following topics:

Terminology

Both the x480 and x880 are the first scalable compute nodes within the Lenovo Flex System family. Each can start as a single compute node 2-socket system.

Scaled complexes

For the x480 compute node, a scalability connector can connect two 2-socket x480 compute nodes, linking them by their QuickPath Interconnect (QPI). A 4-socket 2-node x480 is referred to as a *4-socket x480 complex*.

For the x880, a scalability connector can connect either two 2-socket x880 compute nodes or four 2-socket x880 compute nodes to form either a 4-socket or 8-socket *x880 complex*.

When either an x480 or x880 consists of only a single node (2-socket), it is referred to as a *2-socket system*.

Partitioning options

Both the x480 and the x880 can be logically partitioned. This is where the IMM in the primary node of the complex can logically partition a subset of a 4-socket x480, 4-socket x880, or 8-socket x880 into separate logical servers. For more information, see 5.5.4, "Partitioning" on page 102.

Also, a partition can be made that comprises all nodes within an x480 or x880 complex, otherwise known as a *single-system partition*.

Partitioning occurs at a node level, and partitions control only the physical resources within the nodes in that partition. Partitioning is distinct from more granular I/O virtualization, where a hypervisor (for example, Hyper-V, VMware, or KVM) is used to create virtual machines.

5.5.2 Scalability

The X6 Compute Node complex can scale up to 4-socket or 8-socket configurations, depending upon the processor installed in the nodes:

- The x280 X6 Compute Node with Intel Xeon E7-2800 v2 family processors supports only 2-socket configurations. Each socket must have identical processors installed.
- The x480 X6 Compute Node with Intel Xeon E7-4800 v2 family processors supports 2-socket and 4-socket configurations (4-sockets in a dual-node complex). All processors in the 2-socket or 4-socket configurations must be identical. All other features, such as memory or I/O adapters, can differ.
- The x880 X6 Compute Node with Intel Xeon E7-8800 v2 family processors supports 2-socket, 4-socket (dual-node complex), and 8-socket configurations (*quad-node complex*). All processors in the 2-socket, 4-socket, or 8-socket configurations must be identical. All other features, such as memory or I/O adapters, can differ.

Scaled X6 Compute Node complexes are connected through an interconnecting system located on the front of each X6 Compute Node. The interconnect has all of the main buses (such as QPI), along with the sideband signals, that are required for proper operation.

There are two scalability connector assemblies for the X6 system: A Dual Node (4-socket) Scalability Connector and a Quad Node (8-Socket) Scalability Connector.

The Scalability Connector assemblies for X6 Compute Nodes are listed in 4.5, “Scalability” on page 51. Table 5-4 shows processor, memory, drive and I/O adapter count for each scalable system.

Note: The x480 compute node can be a 2-socket or 4-socket complex. An x880 complex can be a 2-socket, 4-socket, or 8-socket configuration. There is no Scalability Connector that can be used to create a 6-socket complex.

Table 5-4 X6 complex CPU, memory, storage, and I/O comparison

Compute node	Processor sockets	Max DIMM slots	Max drive bays	Max I/O adapter slots
x280 Xeon E7-2800 v2 family				
x280 2-socket	2	48	2	4
x480 Xeon E7-4800 v2 family				
x480 2-socket	2	48	2	4
x480 4-socket	4	96	4	8
x880 Xeon E7-8800 v2 family				
x880 2-socket	2	48	2	4
x880 4-socket	4	96	4	8
x880 8-socket	8	192	8	16

5.5.3 Scalability considerations

When planning to upgrade to a multi-node scalable complex, take these requirements into consideration:

- Additional space within the hosting Flex System chassis

  Secondary nodes in a multi-node complex need to be placed in ascending sequential node bays within their hosted chassis. The design of the 4-socket and 8-socket scalable connectors means that no node bay gaps are allowed. This might mean relocating your primary node farther down in the chassis, relocating nodes currently occupying expansion space, or both.

- Additional power and cooling

  Each additional X6 Compute Node in the complex draws more power from the Flex chassis and requires more cooling. Take this into account in your chassis power policy and power allocation. Table 2-2 on page 17 list power requirements for the X6 Compute Node. For more information on Flex System power allocation and power policies, see "Setting the CMM power management policies" in the Lenovo Flex System Information Center:

  http://ibm.co/1pOT9r7

- Additional I/O module connectivity

  With additional compute nodes installed in the chassis, additional internal ports in all I/O modules must also be activated, which may require additional port licenses. For I/O module requirements, see *Flex System Products and Technology*, SG24-8255:

  http://lenovopress.com/sg248255

- Downtime

  To upgrade to a multi-node complex, you need to allow for downtime. The X6 needs all nodes switched to standby mode while installing the scalable connector.

- Firmware level

  Ensure that the Flex System chassis is at an interoperable firmware level to support the X6 Compute Nodes that you are installing. For details, see 5.5.7, "Partition firmware updates" on page 108.

5.5.4 Partitioning

The Lenovo Flex System X6 Compute Node supports Flex Node Partitioning, allowing a multi-node complex to be sub-divided into smaller physical systems that can operate independently.

Supported partitioning configurations are:

- Single-node configuration (x280 X6, x480 X6 or x880 X6):
  - One 2-socket partition
- Dual-node configurations (x480 X6 or x880 X6):
  - One 4-socket partition
  - Two 2-socket partitions

- Quad-node configuration (x880 X6 only):
  - One 8-socket partition
  - Two 4-socket partitions
  - One 4-socket partition and two 2-socket partitions (the two 2-socket partitions must be adjacent)
  - Four 2-socket partitions

Note: Any 6-socket partitions are not configurable.

Consider the following information when using partitioning a multi-node complex:

- By default, a single partition will be created from all compute nodes in a new X6 complex. If you later add nodes, those will be automatically added to the single partition by default.
- Partitioning is performed at the node level. Single socket (half-node) partitions cannot be created.
- All nodes boot simultaneously so that they are statically configured at the same time. A multi-node complex is not capable of dynamically switching configurations, so all partitions in the complex must be shut down to alter partition configuration.
- All nodes should have the same firmware revisions applied for both UEFI and IMM. For more information, see 5.5.7, “Partition firmware updates” on page 108.
- There are configurations for a multi-node complex where the nodes of a partition need to be in adjacent slots. In an 8-socket x880 X6 complex, nodes 1 and 2 and 3 and 4 are adjacent nodes.
- After initial insertion of a compute node, the system will maintain its previous partition status even after removing a complex, moving compute nodes, flashing firmware, or changing the chassis.
- The node with the lowest node bay ID will be the primary node in the complex, with all other nodes being subordinate. Subordinate nodes receive the configuration information from the primary node through the IMM.
- The primary node disseminates the information to secondary nodes for partitioning boot.
- Single-node partitions that are part of a multi-node complex need to be configured as a single-node partition. Do not leave single-node partitions in the Unassigned Nodes pool. This allows single-node partitions to be able to power up in the event of a partition error condition within the complex.
- After they are created, partitions cannot be expanded (as of the time of writing this chapter). Therefore, a 4-socket partition cannot be expanded to an 8-socket partition, for example. If you need to increase a partition size, you must remove and re-create the partition.
- Partition data remains consistent through reboots, shutdowns, and firmware upgrades.
- Nodes in a partition can have different amounts of memory.

The x480 and x880 compute nodes have a front LED panel to show scaled complex status. For more information on the LED panel, see 4.5, “Scalability” on page 51.

5.5.5 Partitioning by using the IMM web interface

As stated in 5.5.4, “Partitioning” on page 102, a multi-node X6 complex cannot be dynamically partitioned (while the system is powered on). The nodes in the complex that you want to modify must be powered off from the IMM or the CMM.

Before you configure the compute nodes that are part of a multi-node complex, you must partition the multi-node complex by completing these steps:

1. From the IMM web interface, click **Server Management** → **Scalable Complex**, as shown in Figure 5-37.

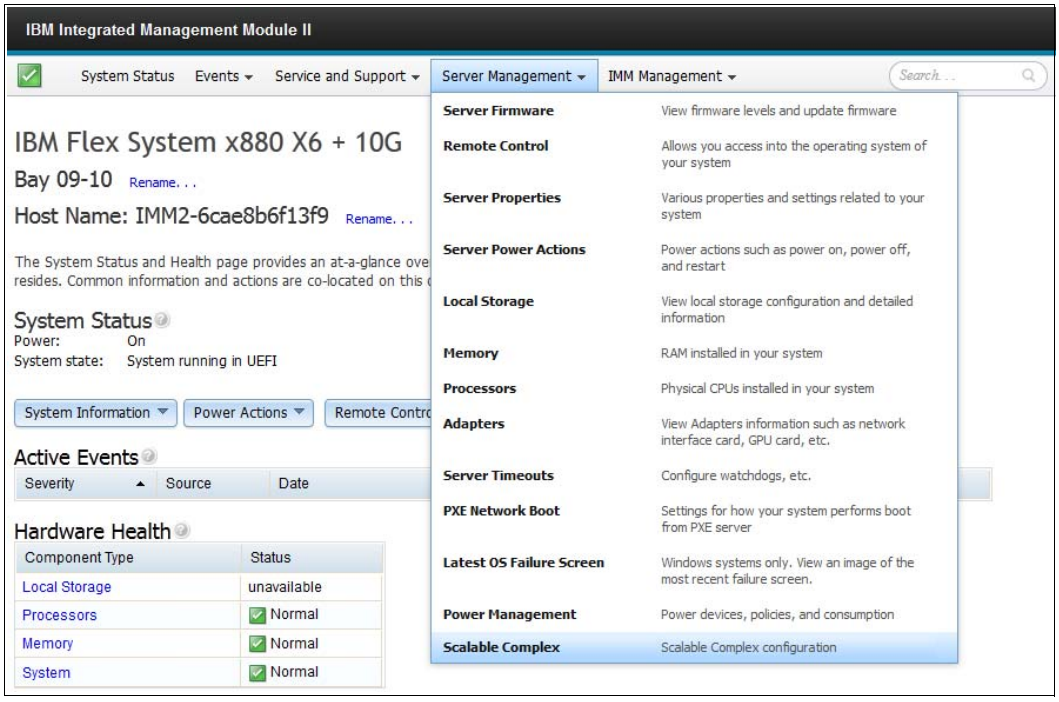

Figure 5-37 Scalable complex server management in IMM web interface

2. From the Scalable Complex window, select the Unassigned Nodes that you want to form into a logical partition, as shown in Figure 5-38

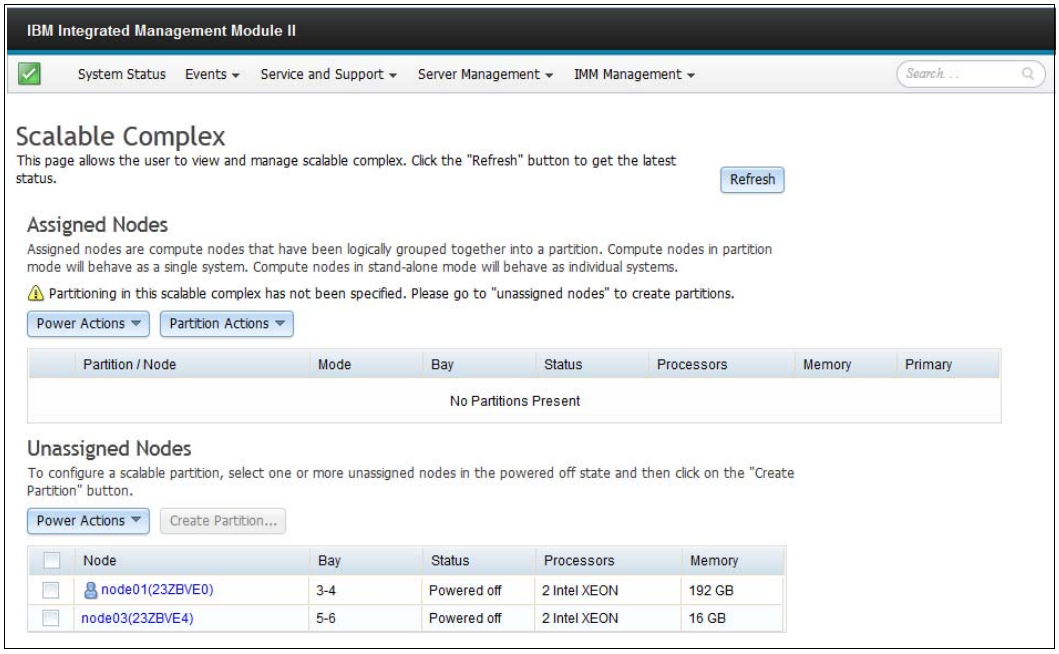

Figure 5-38 IMM web interface Scalable Complex window

3. Select the required nodes for the partition from the Unassigned Nodes table as displayed in Figure 5-39. When all nodes have been selected for the partition, click the **Create Partition** button.

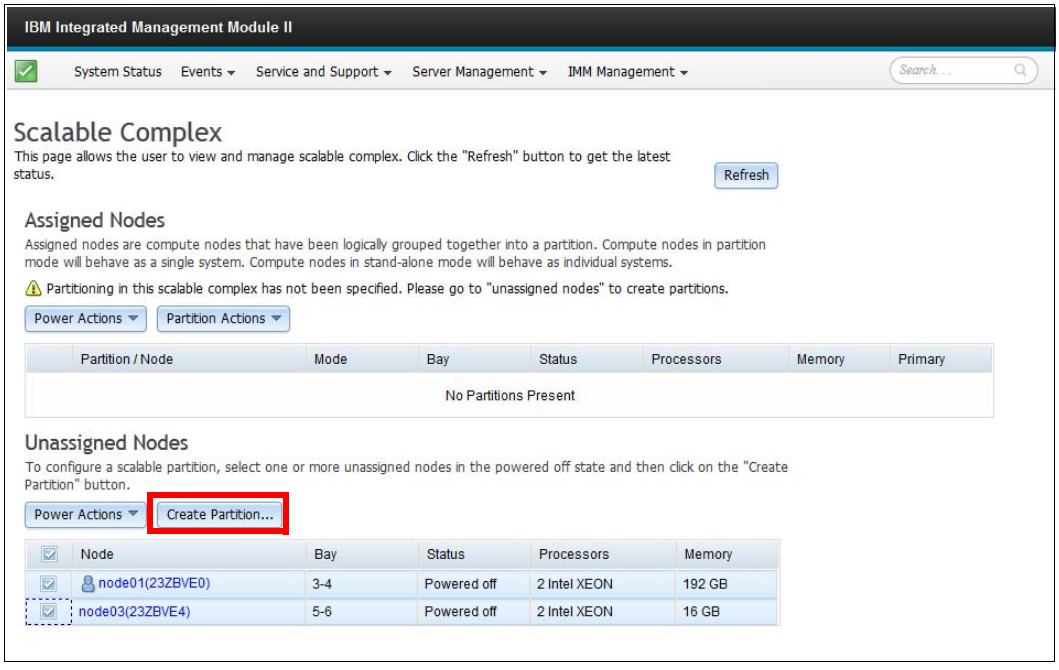

Figure 5-39 Selecting Unassigned Nodes to form X6 partition

4. The Confirm to Create Partition pop-up window provides a summary of which nodes you have selected for the partition and which node will be the Primary Node, as shown in Figure 5-40. Click **Create Partition Now** to complete the process.

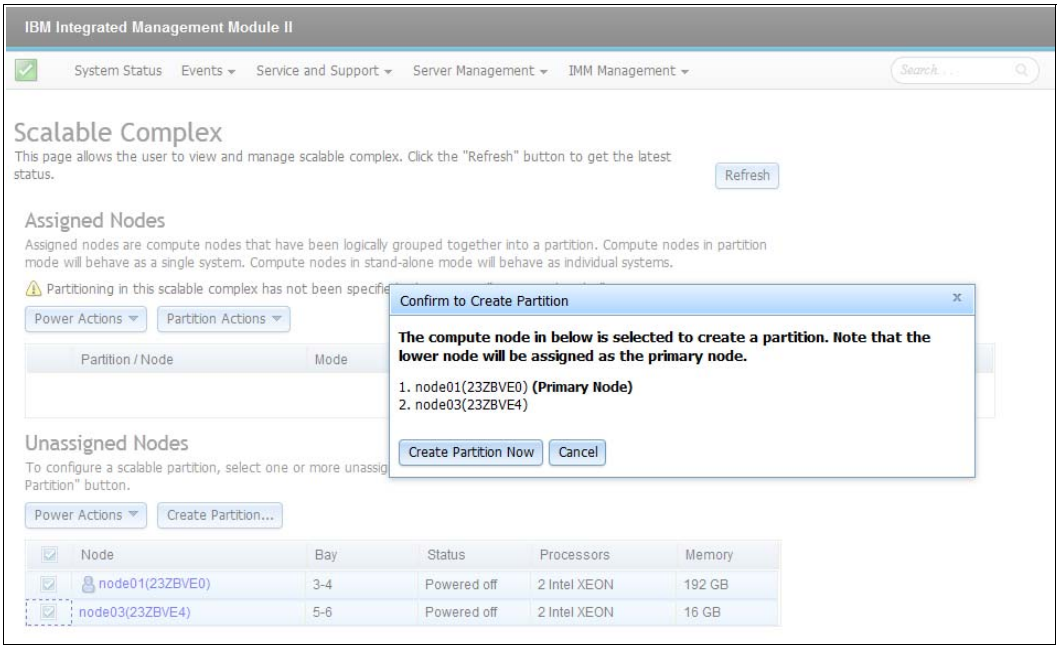

Figure 5-40 Confirm partition creation window

5. Your partition is now ready to be used. To power on the system, select the partition in the Unassigned Nodes table. Then, from the Power Actions drop-down button select **Power On** as shown in Figure 5-41

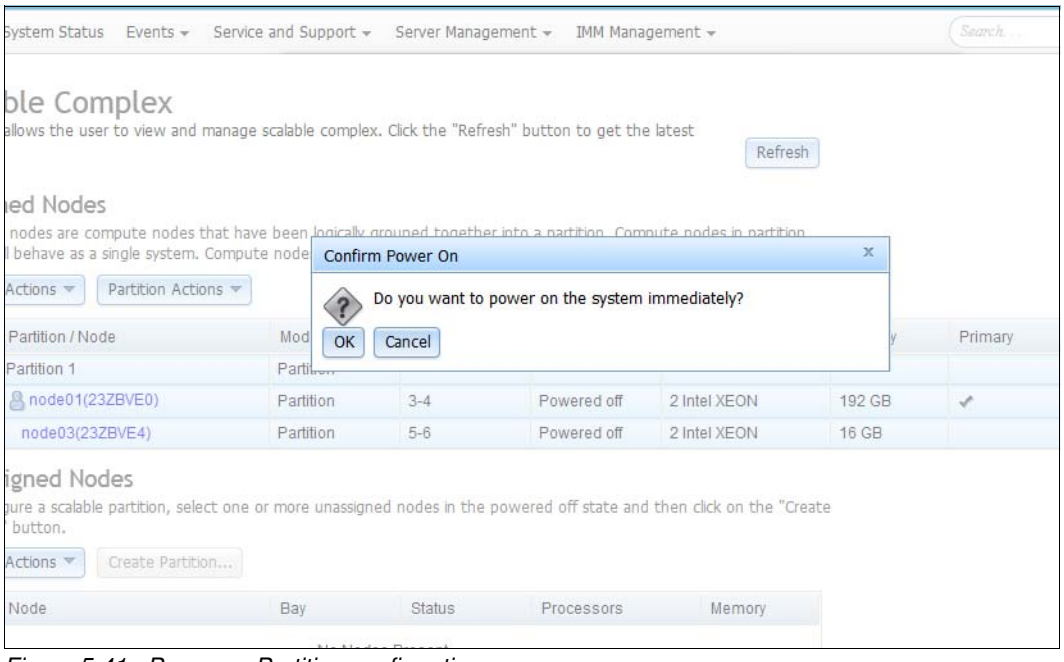

Figure 5-41 Power on Partition confirmation

The partition will now power on and perform a normal POST.

5.5.6 Removing partitions

To remove a partition, it must be shut down first.

1. Shut down the partition from either the operating system for a controlled shutdown or from the CMM or IMM for an immediate shutdown.
2. After the partition is shut down, navigate to the Scalable Complex window from the IMM shown in Figure 5-37 on page 104. Verify that all the nodes in the partition are in a *Powered off* status. Select the partition and from the Partition Actions drop-down tab, and then select Remove Partition, as shown in Figure 5-42.

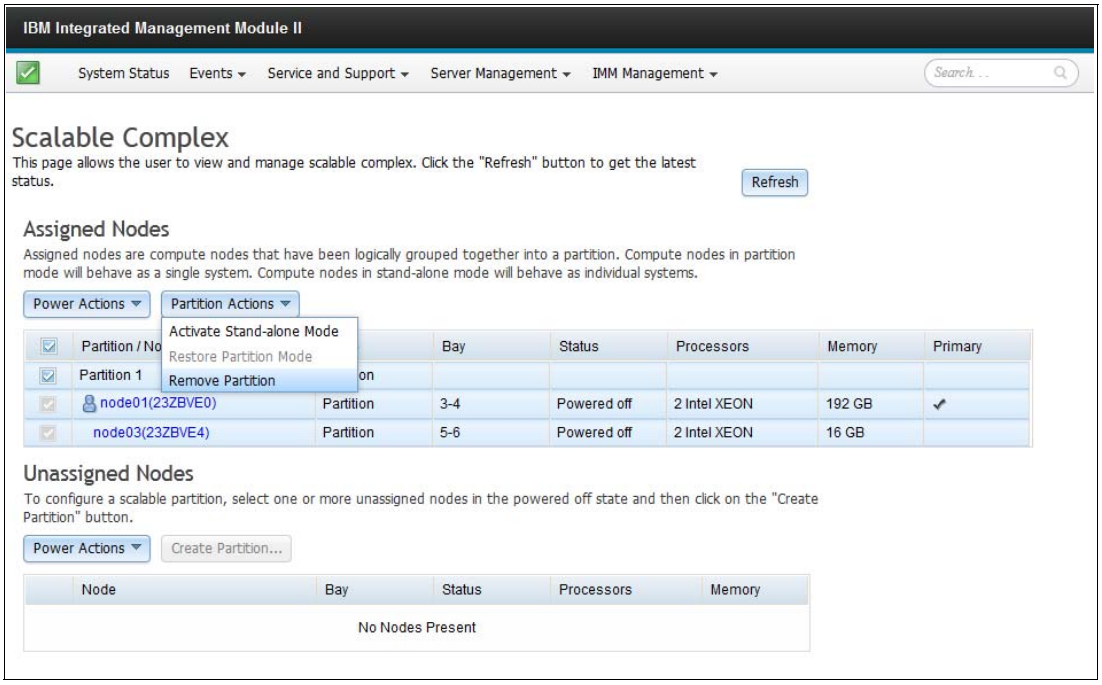

Figure 5-42 Removing a partition from the primary node IMM web interface

All nodes in the partition will move into the Unassigned Nodes table.

Note: If using Drive Sharing ensure that any formed raid arrays have been backed up if required before removing the partition. For more information on Drive sharing please see 5.3.1, “Drive sharing with vertical RAID” on page 94

The Activate Stand-alone Mode switches the partitions to single-node and runs their own UEFI images rather than relying on the primary node within a multi-node partition.

The IMM will not allow partitioning actions to be performed while the selected partition is in a powered on state, as shown in Figure 5-43 on page 108. Any partitioning actions will fail.

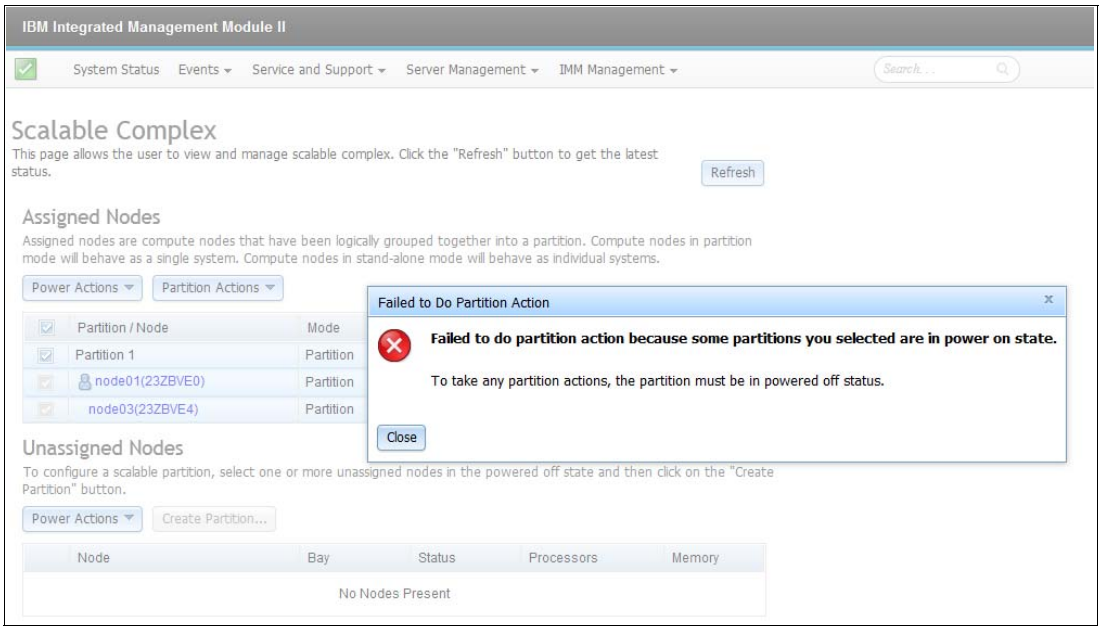

Figure 5-43 Failure of partitioning action because the partition is powered on

5.5.7 Partition firmware updates

Partitions and scaled complexes have firmware considerations that 2-socket systems do not encounter. This is due to the nature of how hardware and partitions are managed.

If you have a single-node system rather than a complex (multi-node), see 6.1.1, “Firmware update tools and preferred practices” on page 114.

In a complex, the first, lowest node bay position node is deemed the *Primary Node* at the *complex* level. This node acts as the arbitration point for information about partitions within the entire complex, as well as actions for partitions such as power control, partition creation, or partition removal.

The IMM in the primary node of the complex communicates with the IMMs in the subordinate nodes, propagating partition information and actions to them. However, the primary node IMM does not have the ability to display information on, or propagate firmware to, the hardware of subordinate nodes in the complex. This means that the process of firmware updates of components such as UEFI, IMM, and DSA in a multi-node complex differs from the process for a single-node system.

Note: It is important here to not confuse the primary node of a *complex* (always the first node of the interconnected physical hardware) with the primary node of a *partition*, which is the first node in a logical partition within the complex. There can be one or more partitions in a complex, so there will be as many primary nodes as there are partitions.

In a multi-node partition, only the UEFI firmware image from the primary node is used upon boot of the partition. Therefore, firmware update tools, such as the Lenovo UpdateXpress System Pack Installer (UXSPI) and the Lenovo Bootable Media Creator™ (BoMC), display only the primary node UEFI. But when the UEFI firmware is upgraded on the primary node, this will be automatically upgraded on the subordinate nodes.

This means that when updating partition firmware the following guidelines should be considered:

- Updating IMM and DSA firmware images for subordinate nodes is possible only through the IMM interface. For information about this method, see “Integrated Management Module” on page 127.
- It is required that IMM firmware for all nodes in the complex be updated to the same level. The CMM will flag an error if there is a mismatch, as shown in Figure 5-44.

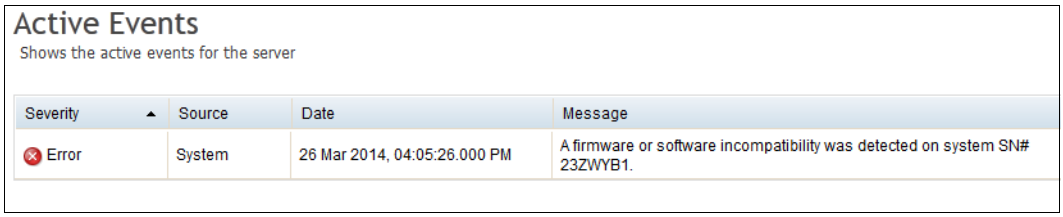

Figure 5-44 CMM and IMM warning of scalable complex IMM firmware mismatch

- Do not restart the IMM of any node in the partition until all IMMs in the partition have the same firmware level in the primary bank, which will become the active level upon IMM restart.
- Avoid using BoMC to update the IMM of the primary node in the partition, because the ToolsCenter will automatically attempt to restart the IMM after upload of the update.
- It is recommended that the UEFI firmware level of all nodes in a complex be kept consistent even if they are subordinate nodes in a partition.

Tip: Due to the complexity of how scalable partitions interact with each other at an IMM level we recommend that all partitions are powered off before updating IMM firmware.

It is still possible and recommended to update drivers and adapter firmware for all nodes in a multi-node partition as you would a single-system partition, using tools such as UXSPI or BoMC.

5.6 Scalable complex installation

To create a scalable complex, first review the planning considerations for physical installation of the compute nodes that are required to form a scalable complex, as discussed in 5.5.3, “Scalability considerations” on page 102.

When installing your x480 or x880 compute nodes, insert them into the Flex System chassis in the lowest compute bay possible, because the complex scales upward in the chassis.

When your first x480 or x880 node is installed into the chassis, the power indicator LED on the compute node will start a fast blink sequence while the IMM is performing its initial load sequence and being discovered by the CMM. The power indicator LED will transition to a slow-blink state when the CMM has discovered the compute node. For more information on the status LEDs on the X6 Compute Nodes, see 3.10, “UEFI” on page 43.

Follow these steps to create a scalable complex:

1. Use the CMM web interface to check whether the compute node has been discovered and is in an OK state. An example is shown in Figure 5-45. If any of the nodes to be interconnected are shaded yellow or red, resolve any issues that have arisen before continuing.

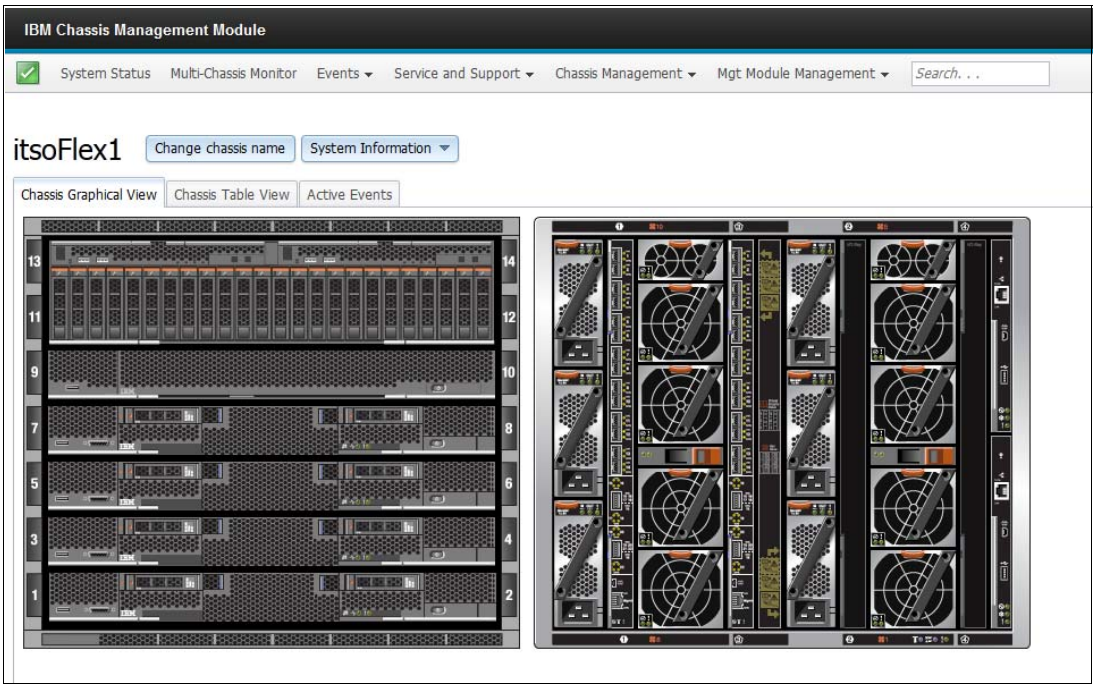

Figure 5-45 Flex Chassis with 4 X6 scalable nodes installed

2. After all nodes for the complex have been installed into the chassis and are in a powered off state, verify that firmware is consistent across all components in all nodes to be attached then afterward attach the scalable connector to the front of the nodes. For more information about scalable connector options, see 4.5, “Scalability” on page 51.
3. After the scalable connector has been installed, verify that the CMM chassis graphical view reflects the change. Figure 5-46 on page 111 shows an example where a 4-socket connector has been installed (which can be used with either the x480 or the x880).

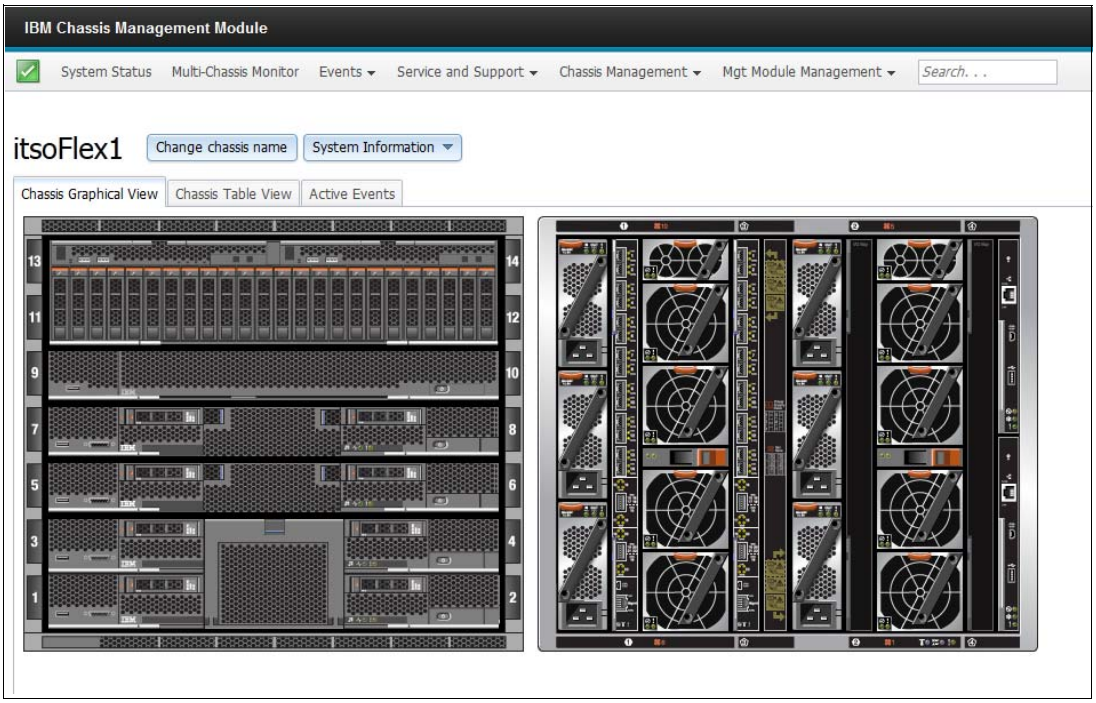

Figure 5-46 4-socket scalable connector attaching nodes in bays 1-4

4. An example of an 8-socket (x880 only) system CMM chassis view is shown in Figure 5-47. If you are upgrading from a 4-socket complex to an 8-socket complex, ensure that all nodes are powered off, remove the 4-socket connector, and then install the 8-socket connector.

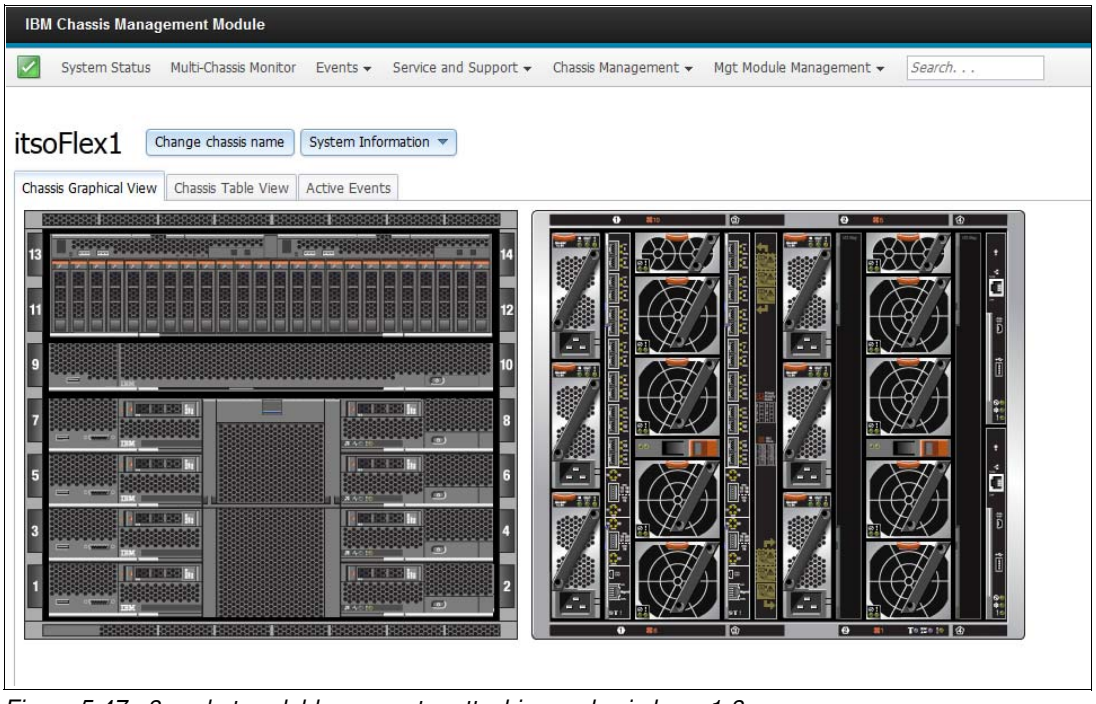

Figure 5-47 8-socket scalable connector attaching nodes in bays 1-8

After the CMM view has updated to your complex configuration, the Scalable Complex option should be available in the IMM. The complex is now ready for partitioning as discussed in 5.5.5, “Partitioning by using the IMM web interface” on page 104.

6

Installation

This chapter describes preparing the Lenovo Flex System X6 Compute Node for production.

We describe firmware updates, how to install an operating system (OS), the operating systems that are currently supported, and provide tips. The chapter includes these two topics:

6.1 Updating firmware

Firmware updates are released to improve functionality, resolve problems that are found, and to support additional hardware and software. During the time in between the system purchase and its first power for use, firmware updates might be released for the server and its components.

When planning to install a system or add options to an existing system (for example, adapters, drives, or memory), firmware levels should be one of the first items to check to ensure that new hardware is compatible with existing hardware or the base system. As new options are added to the supported configuration, some might require a firmware update.

Also, before beginning the installation of any OS or additional hardware, it is best to review the firmware and check whether updates are available. If you are adding hardware after an OS is installed, the suggested first step is to update OS drivers for any possible hardware components. It is best to update drivers before updating firmware.

Lenovo provides many tools to make firmware and driver updates less manual and more streamlined. These tools range from options to update one system to options for deploying updates across multiple systems.

Tools covered in this chapter are available on the Lenovo Support Portal:

http://ibm.com/support

Note: This chapter describes firmware for single-node systems and single-node partitions. For information pertinent to firmware updates for multi-node complexes or multi-node partitions, see 5.5.7, “Partition firmware updates” on page 108.

6.1.1 Firmware update tools and preferred practices

The following firmware update tools that are provided on Lenovo System x machines are described in this section:

- “Fix Central”
- “Lenovo Bootable Media Creator” on page 116
- “Lenovo UpdateXpress System Pack Installer” on page 125
- “Integrated Management Module” on page 127
- “FastSetup” on page 130

These tools streamline the firmware and driver update processes by ensuring that prerequisites and installation orders are met and for more accurate updating with less chance of failure. For more information about the firmware update process for PureFlex and Flex System Compute Nodes, including update methods for firmware and fixes through a Flex System Manager, see the document *Flex System and PureFlex System Firmware Updates Best Practices*, which is available from:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5091991

Follow the relevant release guide when planning any firmware updates to Lenovo Flex System nodes. The guide goes into detail about considerations and how firmware updates can fail if those are not taken into account. If a firmware update fails, review the recovery chapter in that guide. If a firmware update failure cannot be resolved with recommendations in the guide, contact Lenovo Support.

Fix Central

Fix Central is the online source to find all of the available updates for a system type. You can narrow your search by criteria such as a supported component for a particular machine type or the type of OS installed.

Fix Central is available through the Lenovo Support Portal:

http://ibm.com/support/entry/portal/

This support portal is a good starting point for finding all necessary downloads, documentation, upgrade options, and troubleshooting guides for your machine. You can create a login ID or sign in with an existing Lenovo ID to customize a view with any added machine types or previously searched types.

After a machine type is added to your list of searches, you can access Fix Central by using the Download tab at the top of the Support home page, or you can access Fix Central directly and begin your search by machine type:

http://ibm.com/support/fixcentral/

Fix Central can be used to search for, select, and download fixes, updates, drivers, and management or troubleshooting tools for a specific machine type or OS. This feature can be customized by the types of downloads available, such as HTTP, FTP, and Download Director.

You can use Fix Central to find fixes in two ways:

- **Select Product:** Use filters to narrow your selection and then browse for the product.
- **Find Product:** Sarch for a product.

Figure 6-1 on page 116 shows the Fix Central window.

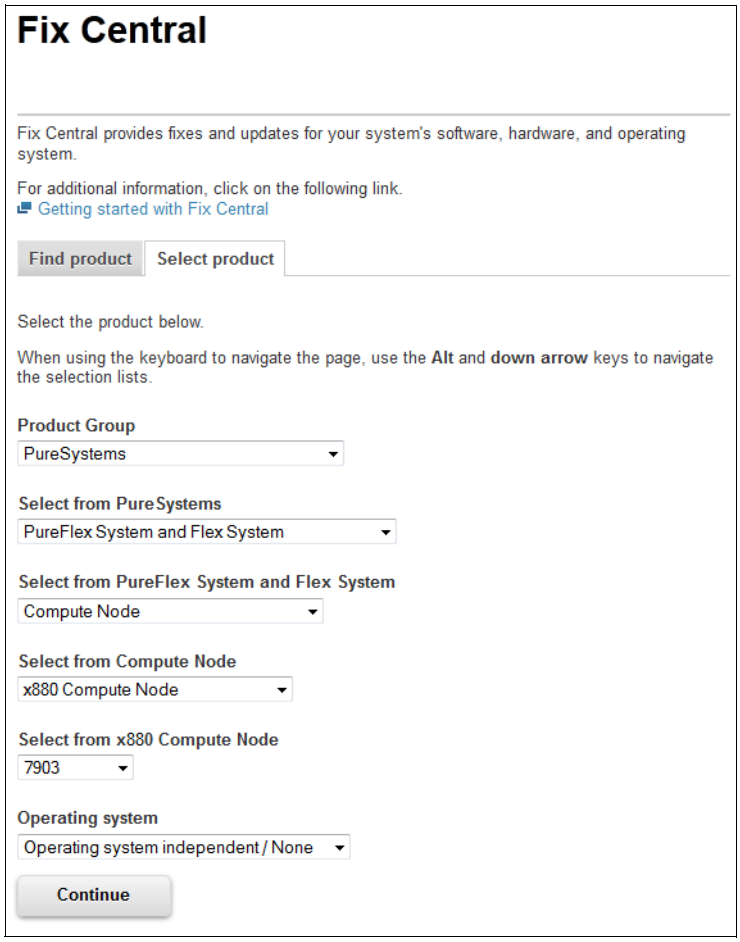

Figure 6-1 Fix Central x880 fix selection

Lenovo Bootable Media Creator

The Lenovo Bootable Media Creator (BoMC) is a tool for the creation of a customized optical media image that is ideal for systems without a functioning OS. This tool is also for operating systems that are not able to run updates or troubleshooting tools. You can get the latest version from the Lenovo Bootable Media Creator (BoMC) web page:

`http://ibm.com/support/entry/portal/docdisplay?lndocid=tool-bomc`

We describe how to use it to deploy updates on a machine-by-machine basis offline.

Consider the following functions when using the BoMC tool:

- When selecting the BoMC tool for an OS, be sure to choose the OS that matches the machine that you are running the BoMC tool on to create the bootable media, rather than the target system where you want to install updates.
- When downloading the BoMC tool, review the `readme` file. This file provides instructions not covered here about how to use the tool.
- When used for the firmware updates, the tool does not update OS drivers. Drivers and firmware levels often need to be kept in sync; therefore, *before* performing firmware updates, ensure that the required drivers are available and updated.

The following examples are taken from the BoMC tool that runs in Microsoft Windows and are used to create update and diagnostic media. The result is to create a customized ISO file either to remote mount through the IMM or to burn to optical disc:

1. Download and run the BoMC tool (administrative access on the local running system will be required). If asked, be sure to read and accept the license agreement.
2. From the Welcome window, you can choose to load a previously saved configuration or to check for the latest version of this tool. Select either one, if applicable. If not, click **Next**. Figure 6-2 shows the Media Purpose window that is after the Welcome window.

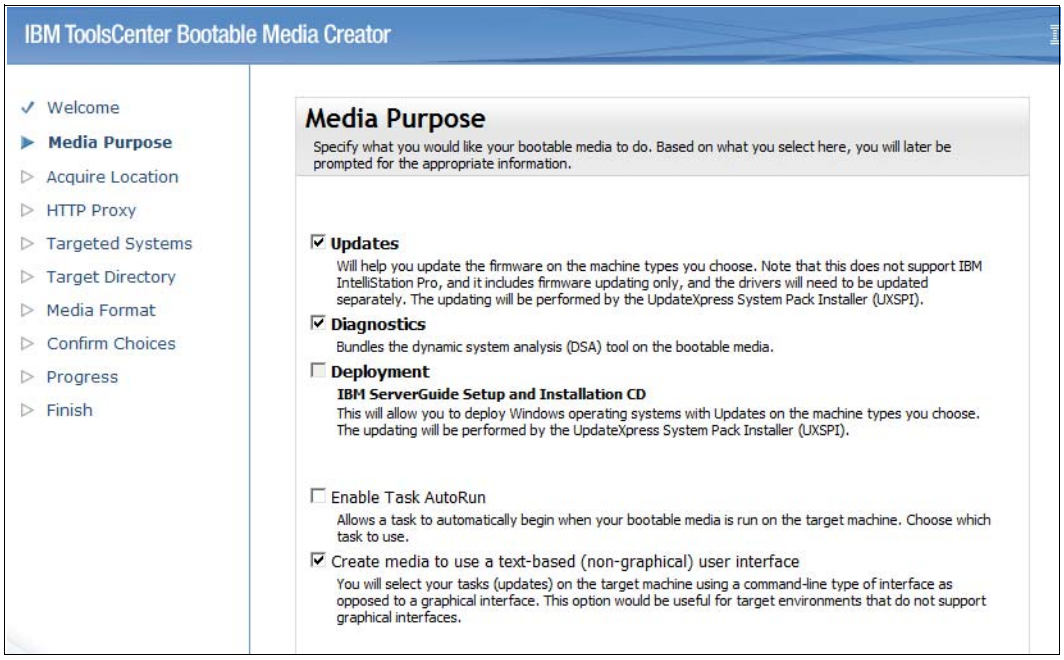

Figure 6-2 BoMC Media Purpose window

By default, the Media Purpose window is set to select *Updates only*. However, this tool can be used for the following tasks:

- Diagnostics: This option bundles in the Dynamic System Analysis™ (DSA) tool.
- Deployment: This option includes the ServerGuide setup and installation media.
- Enable Task AutoRun: This option allows the media to be booted and runs through updates of any firmware files that are contained in the image, without any user intervention.
- Create media to use a text-based (non-graphical) user interface: This option runs only the command-line interface (CLI) version of a BoMC image. It can be useful if updates are run remotely through the IMM remote media mount, where bandwidth is at a premium.

3. With the options selected in Figure 6-2 on page 117, click **Next** to go to the Acquire Location window shown in Figure 6-3.

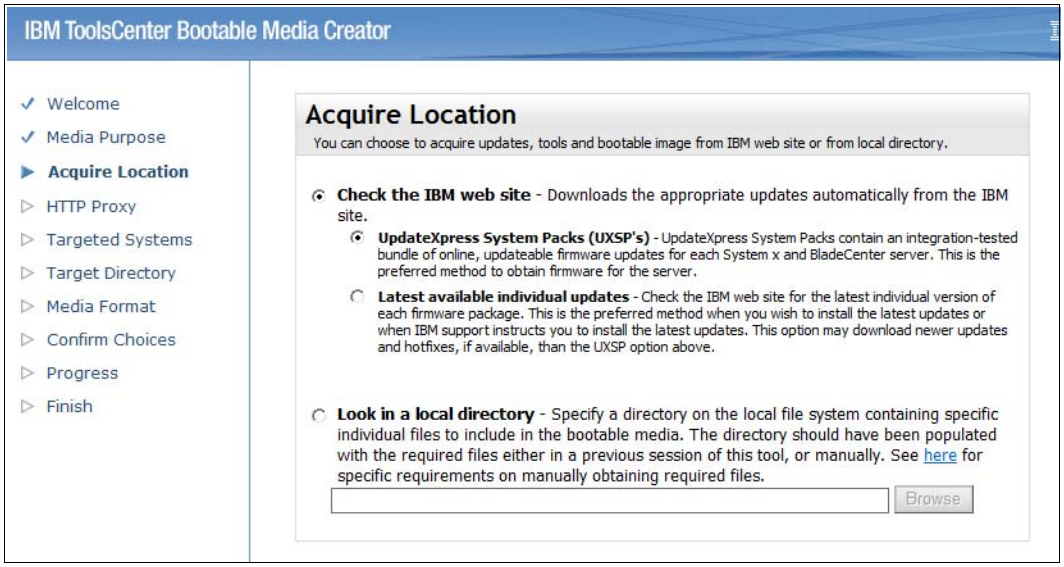

Figure 6-3 BoMC Acquire Location window

4. The Acquire Location function is important to consider when accessing the source of your updates. The following options are available:

   a. Lenovo UpdateXpress System Packs™ (UXSPs): This is the default selection, to receive the UXSP, which is a collection of updates grouped together to ensure compatibility. This default is the suggested selection, but it might not include all of the individual updates that you need.

   b. Latest available individual updates: This option will download and include updates as recent, or newer, than those included in the UXSP option.

   c. Look in a local directory: BoMC can be used to create bootable media in an offline manner (the local running machine has no external network access). This includes a link (hyperlinked) to where to find files to create the bootable environment. This option can also be used to create a more lightweight image with only one or more updates on the image to aid in activity time or when bandwidth to upload the image to the destination remote server is an issue

   After you have made a choice, click **Next**.

5. If you chose either download option in the Acquire Location window, the HTTP Proxy window shown in Figure 6-4 on page 119 opens. Enter HTTP proxy settings if required, or leave the default if you are not using a proxy and direct Internet access is enabled on your local machine.

   After you have made a choice, click **Next**.

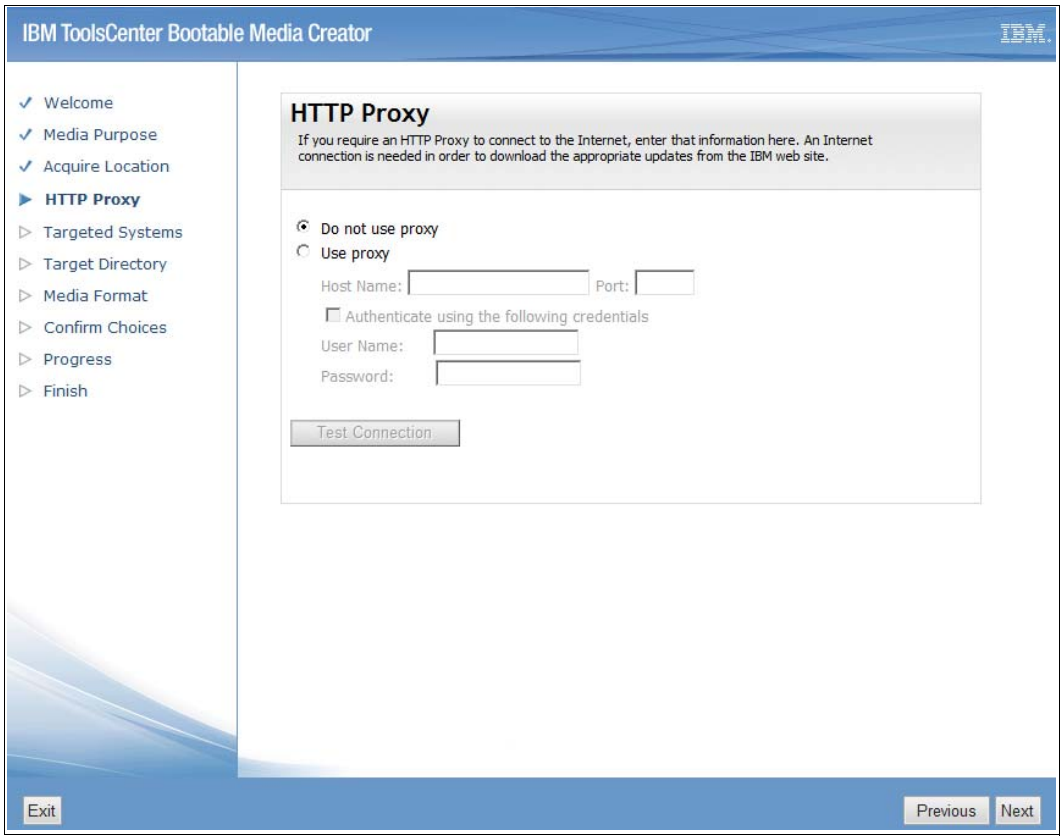

Figure 6-4 BoMC HTTP Proxy window

6. The Targeted Systems window opens next. Choose your targeted system for updates.

   In the example Figure 6-5 on page 120, the Lenovo Flex System x880 system is selected. It is possible to choose more than one machine type so that the bootable media is usable on several systems (for example, the x280, x480, and x880), but that can increase the size of the image file. If the machine type that you need to select is not available in the list, click **Add** and enter the system Product Family and then the Machine Type, and click **OK**. This selection adds the machine to the list to be checked.

   Then, click **Next** to continue.

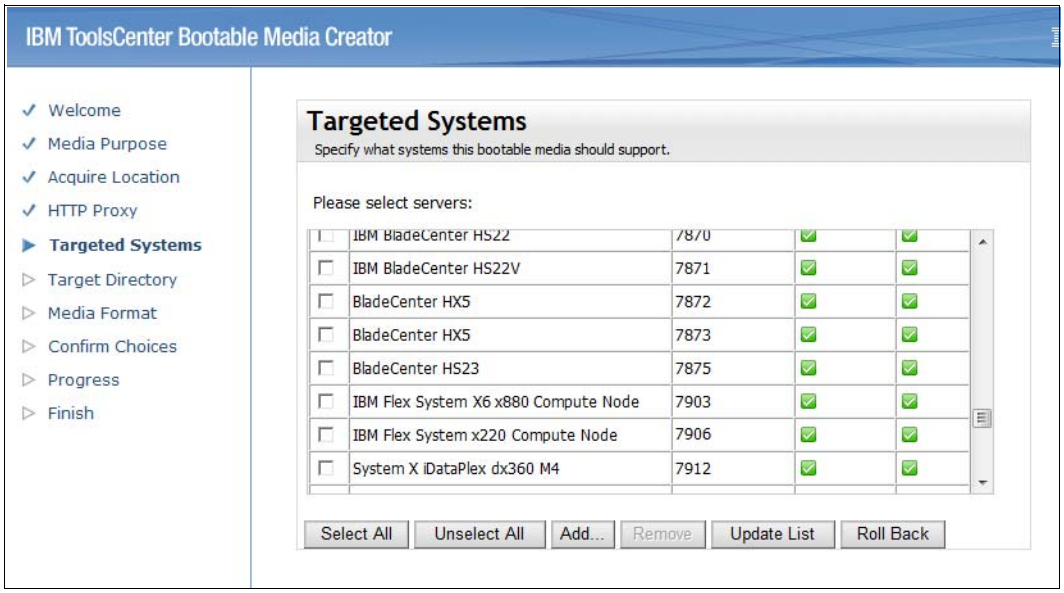

Figure 6-5 Targeted Systems window

7. You will now be asked to specify a Target Directory where files to be downloaded from Fix Central are to be stored. To change the location, click **Browse** to select a different directory.

   Then, click **Next** to open the Media Format window shown in Figure 6-6.

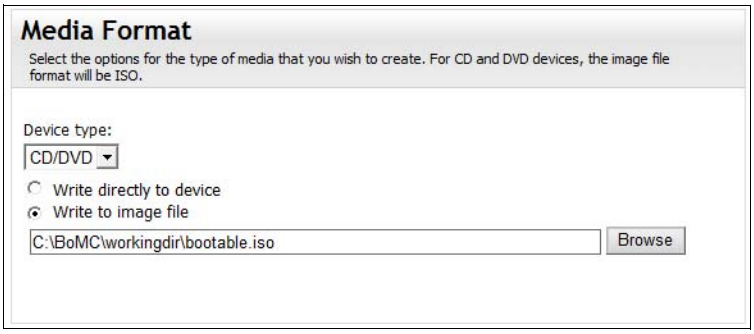

Figure 6-6 BoMC Media Format window

8. In the Media Format window, you can select which type of bootable device to create by using the “Device type” drop-down menu:

   – CD/DVD: Creates an ISO file or allows you to burn directly to an optical device.
   – USB: Formats an attached USB drive and creates a bootable environment on it.
   – PXE: Creates a network-bootable media format to deploy updates over the network.

   Note: By default, Flex Compute Nodes have no optical drive unless an external drive is attached (which may require a USB-powered hub). Therefore, it is better to use an ISO file through IMM Virtual Media or PXE boot (a DHCP-capable PXE boot server will be required).

   If creating an ISO image file, you can choose the folder in which to write the file. BoMC defaults to the directory that u selected previously in the Target Directory window.

   Then, click **Next** to continue.

9. The optional Unattended Mode function helps you configure your created image to run in an unattended mode. This mode uploads the log files to your TFTP server, FTP server, network file share (NFS or Samba), or USB drive. This feature shuts down your clients after the firmware update process. If you want to upload the log files to the TFTP, FTP server, NFS or Samba server, ensure that the directory is created and anonymous access is granted. The default is *not* to use Unattended Mode.
10. In the Confirm Choices window shown in Figure 6-7, confirm your choices in BoMC. Clicking **Save** will save the settings for the next time that they are required.

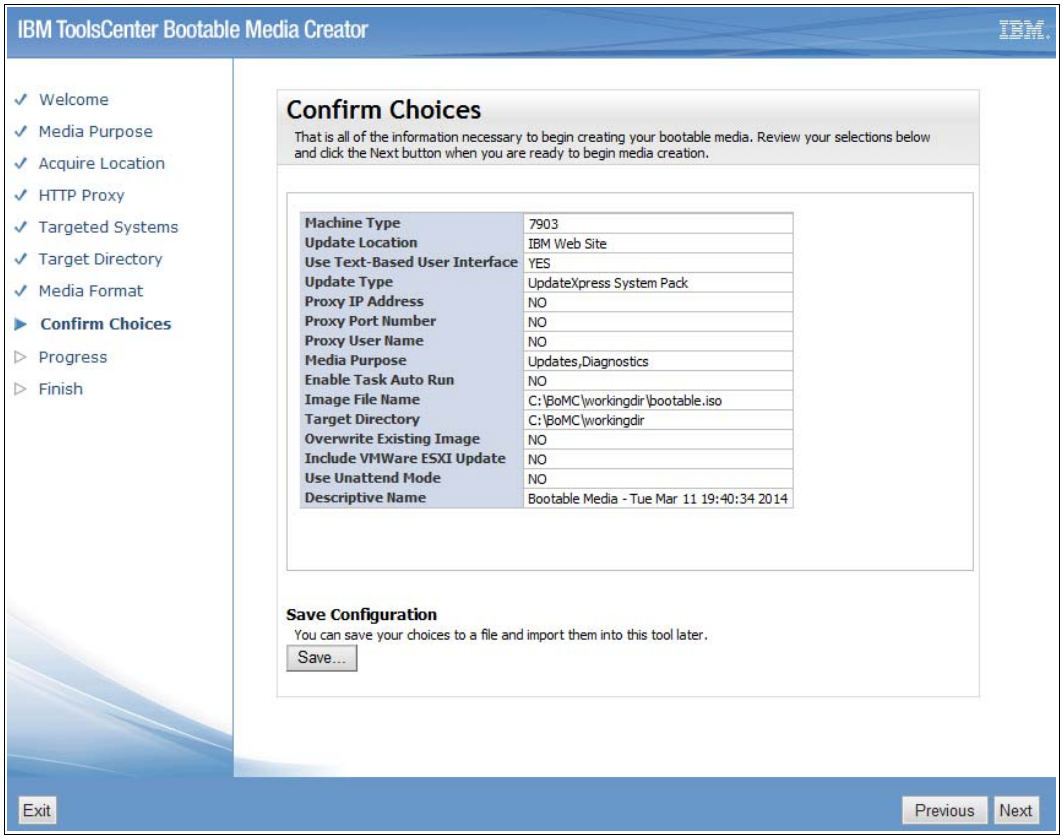

Figure 6-7 BoMC Confirm Choices window

11. The Creation Progress window shown in Figure 6-8 shows the progress while downloading firmware updates and the tools needed to apply them. After all files are downloaded, the bootable media is created. When this step is complete, click **Next**, and then click **Finish**.

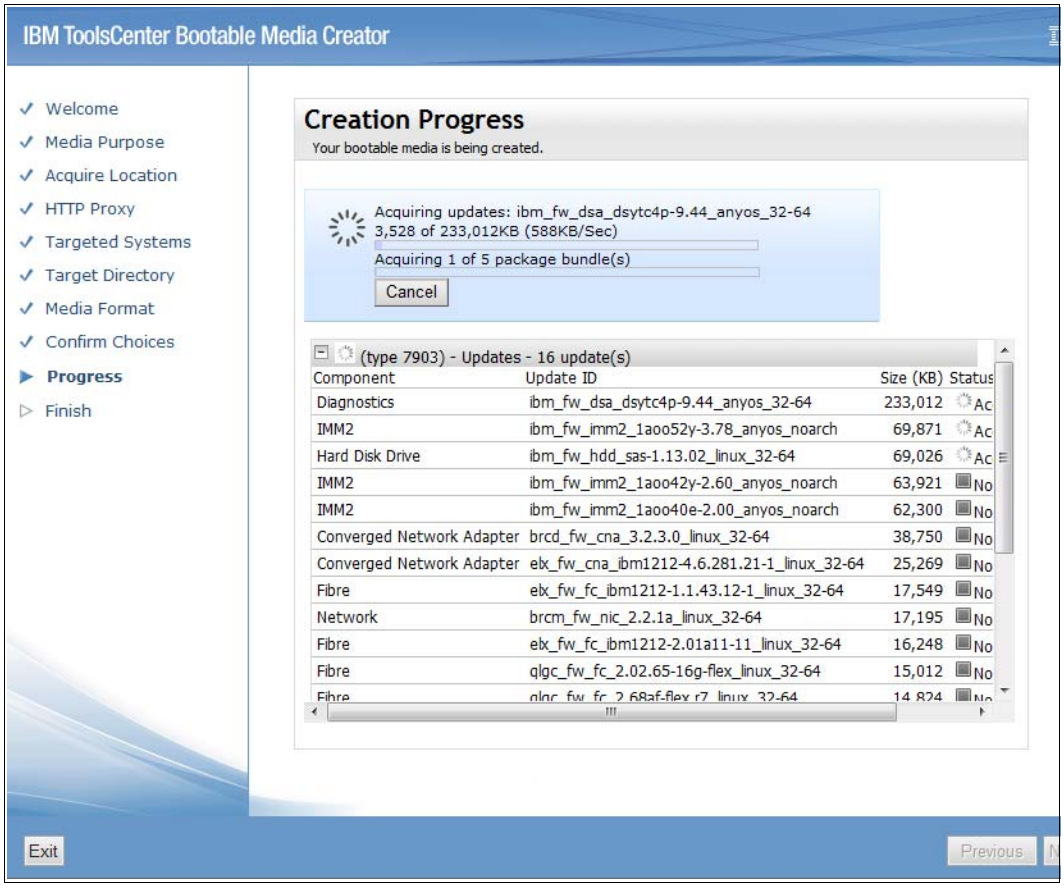

Figure 6-8 BoMC Creation Progress window

In the case of an ISO file, after the bootable media is created, the image can be burned to optical disc. However, it is recommended to use the Remote Control function of the IMM. The media can then be remotely mounted by using the Virtual Media function. Remotely booting the machine is then possible.

Updating firmware with BoMC

1. To start running the tool to perform updates, reboot the server with the mounted media (CD, DVD, or IMM Virtual Media). This reboot launches a CLI or graphical user interface (GUI), as selected in the BoMC image build, to the Lenovo ToolsCenter Customized Media (unless Unattended Mode or Task AutoRun was configured). Figure 6-9 on page 123 shows the Lenovo ToolsCenter Customized Media menu.

```
Welcome to ToolsCenter

IBM System x ToolsCenter is collection server management stand alone tools to help Deploy, Configure, Update and Diagnose your m
achines. This wizard will help you manage your system more efficiently.

Discriptive name:
 Bootable Media X6 x880 updates - Wed Mar 12 14:40:34 2014


ToolsCenter have following function:
 <1> Updates

Enter the item number to Updates| ('q' to quit ToolsCenter):
```

Figure 6-9 ToolsCenter initial menu

2. Select option **1** for updates. Depending upon the tasks that you selected in BoMC, you will see an additional option for Diagnostics.

   The process will now load the media to determine what updates are available and then inventory the server to cross-reference what updates are applicable to the server.

   Figure 6-10 shows an example of the window that lists updates that are available.

```
UpdateXpress System Pack Installer Version 9.51.06        Machine Type 7903

Press any key to cancel countdown timer.                                                    Automatic upd
Updates automatically start when countdown reaches 0.
-------------------------------------------------------------------------------------------------------


  1 [*]Integrated Management Module 2 (IMM2) Update
         Category          :IMM2
         Severity          :Critical
         Reboot            :Not Required
         Update ID         :ibm_fw_imm2_1aoo52y-3.78_anyos_noarch
         Requisites        :None
         New Version       :3.78 (1AOO52Y)
         Installed Version :2.00 (1AOO40E)

  2 [*]IBM x880 UEFI Flash Update
         Category          :UEFI
         Severity          :Suggested
         Reboot            :Reboot Required to take effect
         Update ID         :ibm_fw_uefi_b2e136u-1.41_anyos_32-64
         Requisites        :None
         New Version       :1.41 (B2E136U)
         Installed Version :1.21 (B2E122D)

  3 [*]Emulex HBA (LPe1205/LPe1200x) Firmware Update for Linux - 2.01a11-5.12a15 - Release IBM1212
         Category          :Fibre
         Severity          :Suggested
         Reboot            :Reboot Required to take effect
         Update ID         :elx_fw_fc_ibm1212-2.01a11-11_linux_32-64
         Requisites        :None

         Update            :95Y2377 Firmware (#1)
         New Version       :2.01A11
         Installed Version :2.01A9

         Update            :95Y2377 Firmware (#2)
         New Version       :2.01A11
         Installed Version :2.01A9

         Update            :95Y2377 Bios (#1)


-------------------------------------------------------------------------------------------------------

List View       Select All       Apply All Selected       Quit       Help
```

Figure 6-10 UXSP firmware updates application window

3. When the list of suggested updates is compiled, you have *60 seconds* to stop the tool from continuing and applying all possible *critical* and *recommended* updates found. The *suggested* updates option (for example, a minor revision of HBA firmware) will be left clear. As soon as any key is pressed, the countdown timer is halted and you can review and either select or clear updates.

   The update process completes as the prerequisites require. When the updates are complete, a summary window will be displayed showing which updates were applied as shown in Figure 6-11.

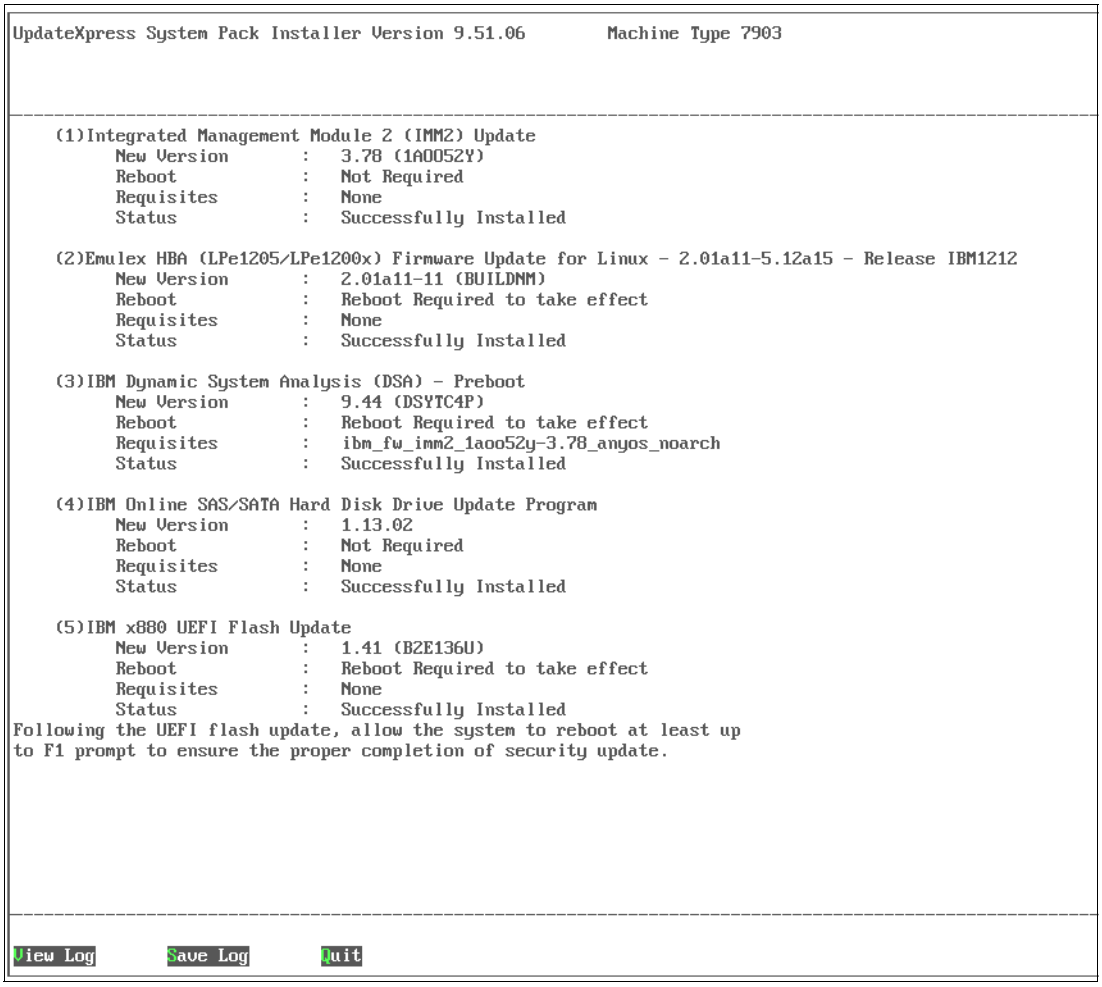

Figure 6-11 UXSP Update summary page

4. The machine can now be rebooted by selecting **Q** to quit, which exits to the main ToolsCenter menu.

   Then, click **Q** again to quit the ToolsCenter, where media will be ejected or unmounted and the server will be rebooted.

> **UEFI firmware update:** If a UEFI firmware update was included as part of the updates, during the first POST after quitting ToolsCenter a message stating '**Warning: Secure Region update in Progress, Do Not Power Off System**' will be presented. After this message the system will then perform a 2nd POST and updated firmware levels will be reflected both in the IMM Firmware window and on the CMM node Firmware tab. It is imperative that power is not interrupted during this process or any user-initiate reboot is performed, as either can cause erratic system behavior.

Lenovo UpdateXpress System Pack Installer

The Lenovo UpdateXpress System Pack installer (UXSPI) is a tool that can be used to update a server directly on the booted OS. With UXSPI, you can also create a repository of firmware and driver updates to be applied later.

UXSPI uses the standard HTTP (port 80) and HTTPS (port 443) protocols to get the updates from Lenovo. Your firewall must allow these ports. UXSPI is supported on Microsoft Windows, Linux, and VMware operating systems and is supported on both 32-bit and 64-bit operating systems.

The *Lenovo UpdateXpress System Pack Install User's Guide* provides a detailed list of supported operating systems. It is available by using the Publications link on the Support Portal web page:

http://ibm.com/support/entry/portal/docdisplay?lndocid=SERV-XPRESS

Follow these steps to use UXSPI:

1. Download and then start the UXSPI setup utility from the website.
2. Accept the license agreement.
3. After the Welcome window shown in Figure 6-12. opens, click **Next**.

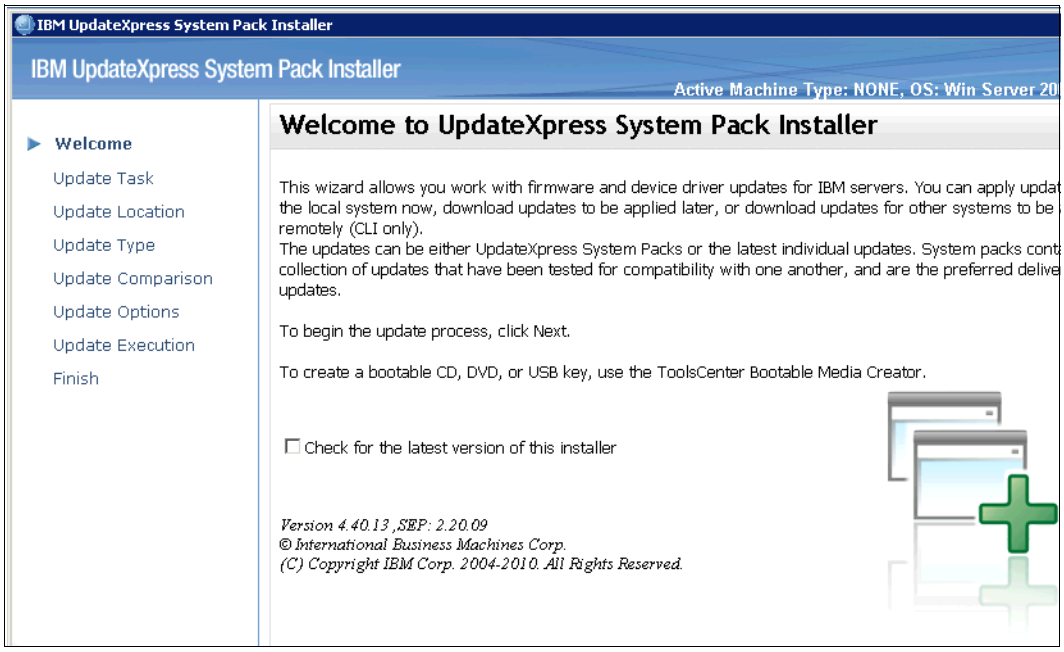

Figure 6-12 UXSPI Welcome window

4. Accept the **Update the local machine** default, as shown in Figure 6-13, and click **Next**.

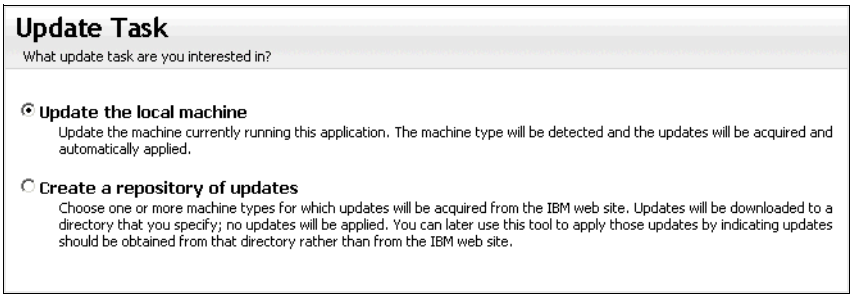

Figure 6-13 UXSPI Update Task panel

5. Accept the **Check the IBM website** default, as shown in Figure 6-14, and then click **Next**.

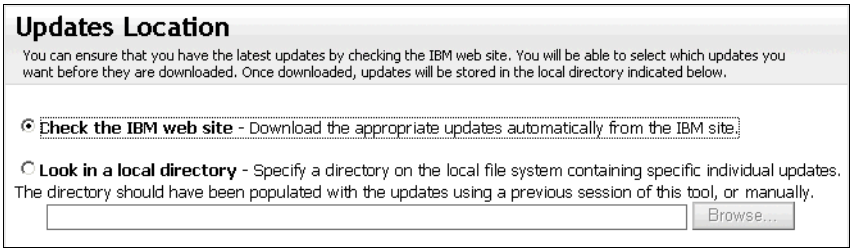

Figure 6-14 UXSPI Updates Location panel

6. It is important to consider when which update type to choose, as shown in Figure 6-15:
   a. **Lenovo UpdateXpress System Packs (UXSPs):** This is the default selection to receive the UXSP, which is a collection of updates grouped together to ensure compatibility. This default is the standard suggested selection, but it means that certain latest individual updates might not be included.
   b. **Latest available individual updates:** This option downloads and includes updates as recent, or newer, than those included in the UXSP option.

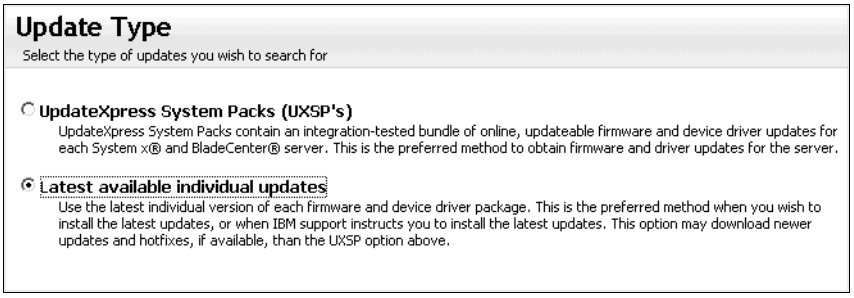

Figure 6-15 UXSPI Update Type panel

7. Now, select the Target Directory where you want to store the files to be downloaded, as shown in Figure 6-16, and then click **Next**.

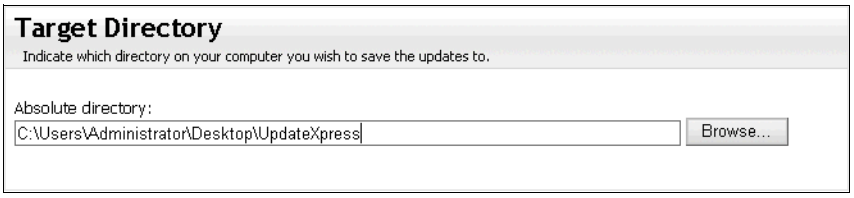

Figure 6-16 UXSPI Target Directory panel

8. Enter the settings for an HTTP proxy server, if necessary, or leave the check box clear, as shown in Figure 6-17. Then, click **Next**.

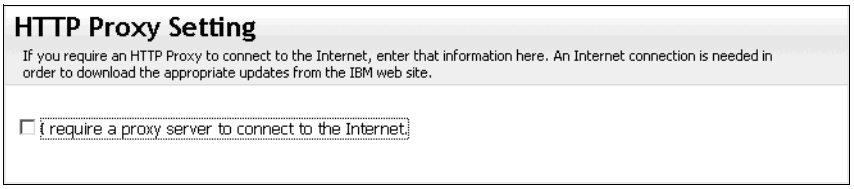

Figure 6-17 UXSPI HTTP Proxy Settings panel

9. A message displays showing that UXSPI acquired the possible updates for the machine, as shown in Figure 6-18. Click **Next**.

Figure 6-18 Successful completion of acquisition message

10. The next message appears showing that the download has completed, as shown in Figure 6-19. Click **Next**.

Figure 6-19 Download process completion message

11. A component overview shows the components that need updating. UXSPI selects the components to update by default. Accept these settings, and click **Next**.
12. When the update is finished, a message displays confirming that the updates have been applied successfully, as shown in Figure 6-20. Click **Next**.

Figure 6-20 Update process completion message

13. Click **Finish** to close UXSPI.
14. Restart the system to complete the update process.

Integrated Management Module

The IMM supports updating of firmware by using either the web interface or the CLI. The Integrated Management Module (IMM) enables you to update DSA, UEFI, and IMM firmware.

Note: Updating server firmware by using the IMM is intended for recovery purposes. The preferred and recommended methods of updating firmware are by using either Lenovo UpdateXpress System Pack Installer (UXSPI), as explained in “Lenovo UpdateXpress System Pack Installer” on page 125, or the Lenovo Bootable Media Creator (BoMC) as explained in “Lenovo Bootable Media Creator” on page 116.

The following section describes the process of updating the IMM, UEFI, and DSA by using the IMM web interface.

1. From the IMM web interface, select **Server Management** → **Server Firmware**, as shown in Figure 6-21.

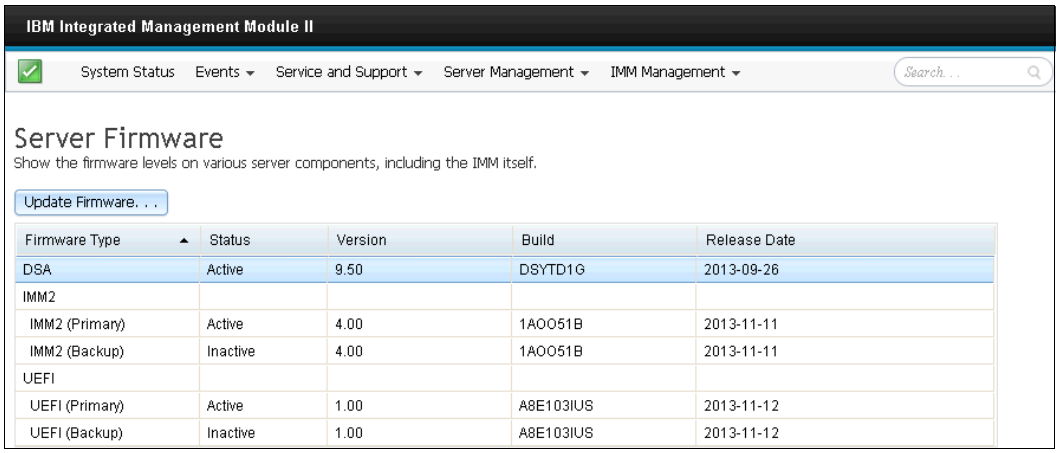

Figure 6-21 The IMM Server Firmware window

2. Click the **Update Firmware** button. The Update Server Firmware wizard shown in Figure 6-22 will open to guide you through the update process.

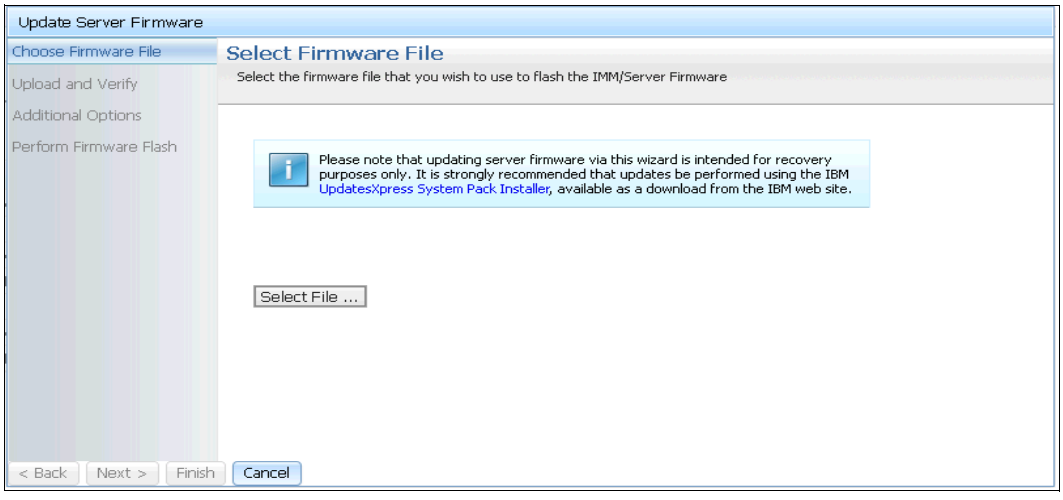

Figure 6-22 IMM Select Firmware File panel in the Update Server Firmware wizard

3. Ensure that you have downloaded the appropriate firmware update from Fix Central or relevant support page. The IMM has no option to download files from IBM. After the file is selected, click **Next.**

   A progress bar will show the file being uploaded to the IMM and verified. Afterward, a panel showing current firmware revision level and the level the uploaded file will be displayed, as shown in Figure 6-23.

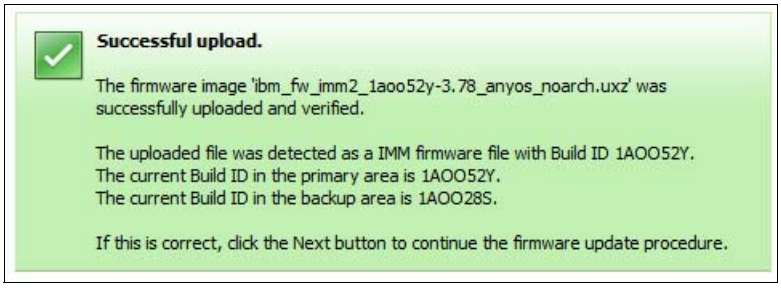

Figure 6-23 IMM Firmware Upload verification information panel

4. After the file is verified to be the level that you require, click **Next** to open the Additional Options panel, where you can choose the action, as shown in Figure 6-24:
   - Action1, the default, is to apply the file to the Primary bank.
   - Action2 can be selected so that both Primary and Backup banks can be updated simultaneously.
5. When you have selected the action, click **Next** to begin the update.

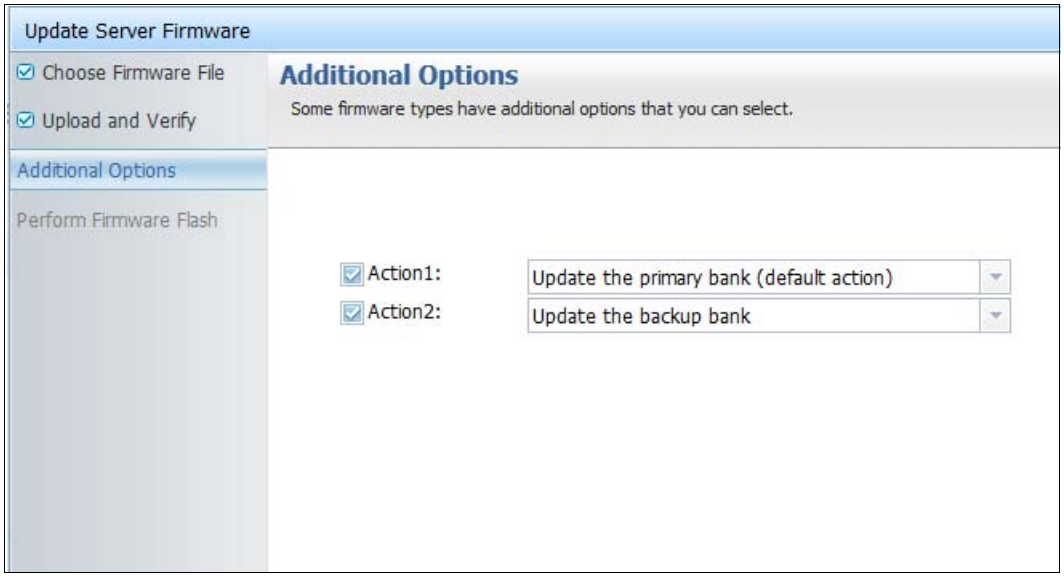

Figure 6-24 IMM Firmware Action panel

6. A progress bar shows firmware update status. When the update is complete, the “Firmware update success” message is displayed, as Figure 6-25 shows.

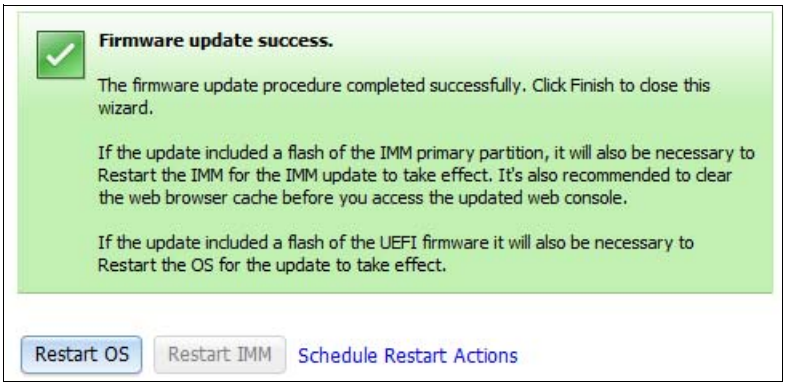

Figure 6-25 “Firmware update success” message

Note: The **Restart OS** or **Restart IMM** buttons perform those actions if clicked when they are applicable. If an update is applied to a Backup bank for the IMM or UEFI, it becomes active immediately, without a restart.

7. Click **Finish** to quit the Firmware Update wizard without rebooting. If a restart of the IMM or a reboot of the OS is required, a message stating they are required so that firmware updates can be activated is displayed, as shown in Figure 6-26.

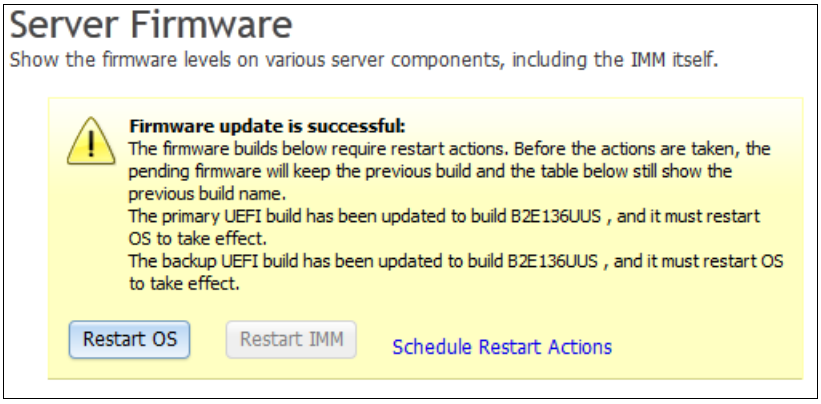

Figure 6-26 Firmware activation pending restart message

FastSetup

FastSetup enables you to create mass deployments of firmware updates and includes features to make the setup process as simple and as fast as possible. FastSetup includes automation templates for mass deployments of firmware, which makes it easy to push configuration settings. It includes integrated help for dialog panels.

FastSetup runs on Windows and performs the following tasks:

- Opens a setup wizard to collect inventory, update firmware, set configuration parameters
- Creates a template of firmware levels and parameters to be deployed to other servers or I/O modules of the same type
- Downloads firmware to a local repository for later use

The system running FastSetup needs to be able to connect to the target systems through the network, but it does not have to be on the same subnet.

Complete these steps to install FastSetup:

1. Download the FastSetup utility from the FastSetup page:

   `http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-FASTSET`

2. Launch the tool from your download location.
3. Select your language, and click **OK**.
4. After the introduction page opens, click **Next** and follow the instructions to install the program. Figure 6-27 shows the installation GUI and steps taken to install.

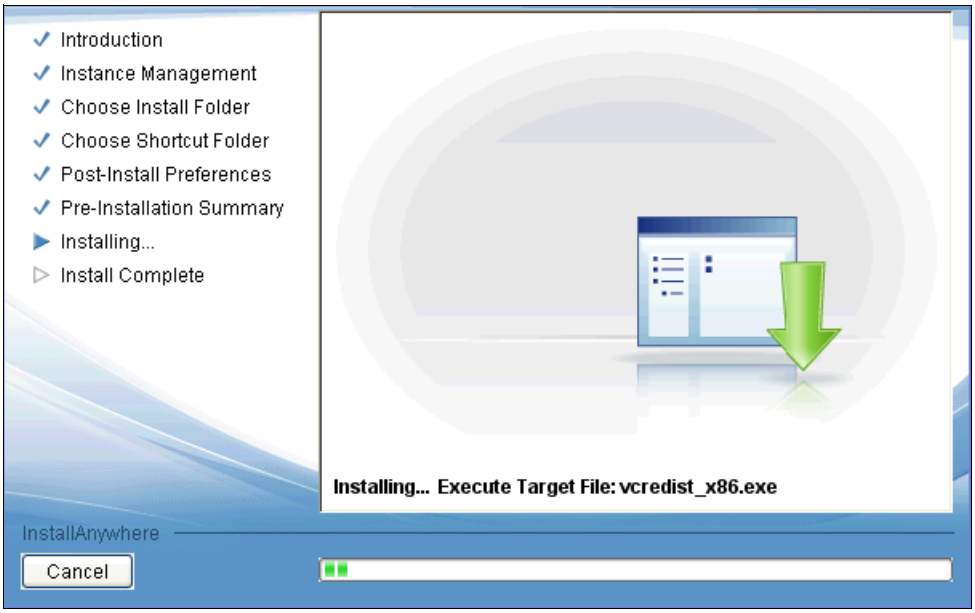

Figure 6-27 FastSetup installation

5. When the installation is complete, you will receive a success message in the Install Complete panel shown in Figure 6-28. Click **Done**.

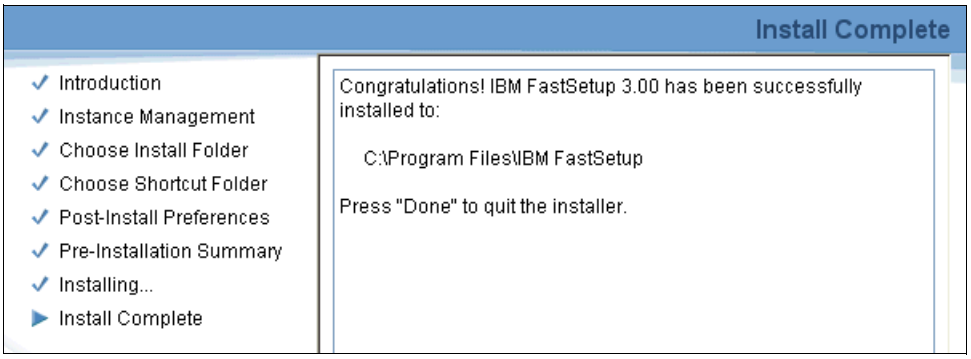

Figure 6-28 Installation Complete message

6. Launch the FastSetup utility from your installation location by clicking on the `FastSetup.exe` file.
7. Select your language and accept the terms and conditions. The Welcome page that Figure 6-29 shows then opens.

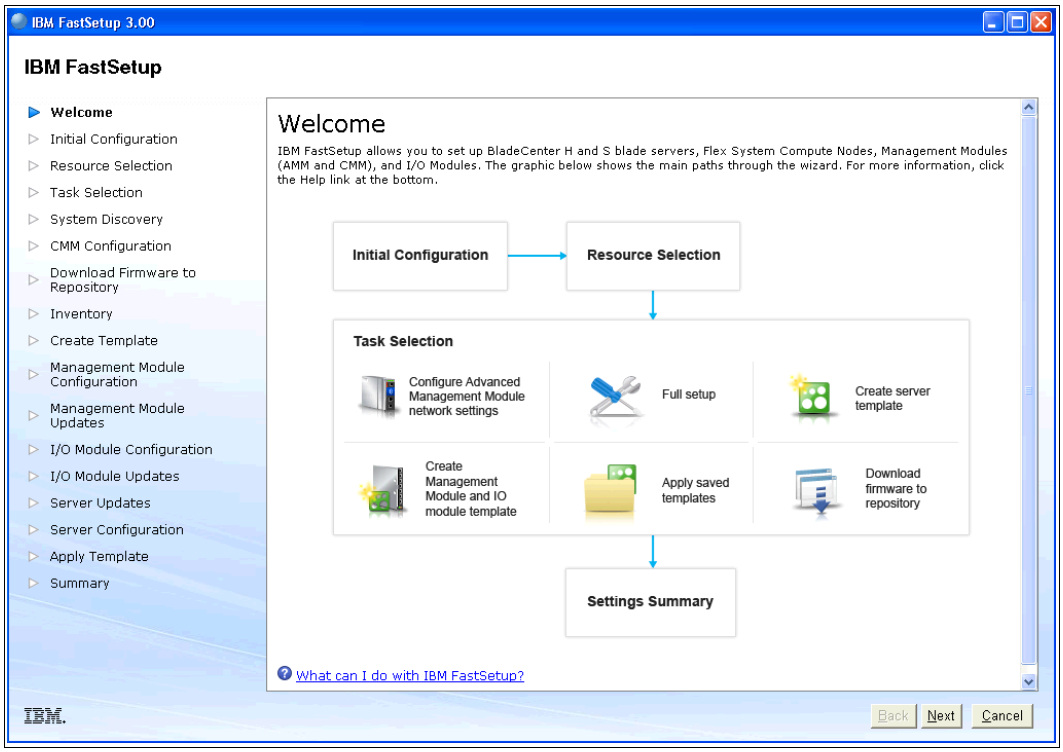

Figure 6-29 FastSetup welcome page

8. To use the tool, follow the procedures in the FastSetup white paper:

   `http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5092092`

6.2 Installing an operating system

The following topics in this section provide an overview of the options that you have when installing an OS on the Lenovo Flex System X6 family of compute nodes:

- 6.2.1, “Supported operating systems” on page 133
- 6.2.2, “Installing without a local optical drive” on page 133
- 6.2.3, “Lenovo ServerGuide” on page 134
- 6.2.4, “Lenovo ServerGuide Scripting Toolkit” on page 140
- 6.2.5, “Use of embedded VMware ESXi” on page 141
- 6.2.6, “Booting from SAN” on page 142

An OS can be installed on the Lenovo x880 with either local or remote media. The x880 does not have any onboard optical device but can use an external USB optical drive or USB media to perform a local installation.

6.2.1 Supported operating systems

The x880 compute node is tested to support many server operating systems that are currently available. This section lists the available options that are supported by Lenovo to run on the x880 compute node. For detailed information about the process used to approve software and hardware for support in Flex System, see the ServerProven web site:

http://ibm.com/systems/info/x86servers/serverproven/compat/us/flexsystems.html

The list of supported operating systems is in 4.17, “Operating systems support” on page 70.

The operating systems listed are supported with the X6 Compute Nodes. When an operating system is listed as supported by Lenovo, it passed rigorous testing. Also, Lenovo worked with the operating system vendor to provide the necessary support for the base node and options.

6.2.2 Installing without a local optical drive

The X6 family of compute nodes do not have a local optical drive. To install an OS you can use any of the following methods, which are described in this section:

- “Integrated Management Module” on page 133
- “Local USB port” on page 133
- “Preboot eXecution Environment” on page 134

Integrated Management Module

The remote control KVM feature is available through the IMM web interface. You must log in to the IMM with a user ID that has Supervisor access. You can also assign an ISO image, USB, or optical device to the server from your local machine by using the IMM Virtual Media function, as covered in 5.1.3, “Accessing the remote control feature in the IMM2” on page 74.

For more information, see the *Integrated Management Module II User’s Guide* avaliable from:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086346

Local USB port

You can use the local USB port to attach a USB device that contains the OS installation files. This can either be a USB flash drive or a USB optical device.

Note: Some CD and DVD USB-powered devices might not be compatible, and an external USB-powered hub might be required for optical media installation.

The following vendor web pages provide details for using a USB key as an installation medium:

- Installing Red Hat Linux from a USB flash drive, in the *Red Hat Enterprise Linux Installation Guide*:

  http://red.ht/1rqBnNz

- "How to create a bootable USB drive to install SLES," in the Novell Knowledgebase:

  http://www.novell.com/support/kb/doc.php?id=3499891

- "Install Microsoft Windows from a USB flash drive," in the Microsoft TechNet Library:

  http://technet.microsoft.com/en-us/library/dn293258.aspx

- "Format a USB flash drive to Boot the ESXi Installation or Upgrade," in the VMware vSphere Documentation Center:

  http://vmw.re/XG7lbu

You can also use the ServerGuide Scripting Toolkit to create a bootable USB flash drive, as explained in 6.2.4, "Lenovo ServerGuide Scripting Toolkit" on page 140.

Preboot eXecution Environment

The Preboot eXecution Environment (PXE) is an environment to boot computers by using a network interface for operating system deployment. All X6 systems support PXE booting.

For example, you can use xCAT software to deploy a broad set of operating systems over the network. For more information, check the xCAT web page on the SourceForge website:

http://sourceforge.net/apps/mediawiki/xcat

6.2.3 Lenovo ServerGuide

Lenovo ServerGuide is an installation assistant for Windows installations that simplifies the process of installing and configuring Lenovo Flex System, Lenovo System x, and Lenovo BladeCenter servers. The wizard guides you through the setup, configuration, and operating system installation.

The ServerGuide can accelerate and simplify the installation of X6 Compute Nodes in the following ways:

- Assists with installing Windows-based operating systems and provides updated device drivers that are based on the detected hardware
- Reduces reboot requirements during hardware configuration and Windows operating system installation so you can get your X6 Compute Node up and running sooner
- Provides a consistent server installation, using Lenovo best practices for installing and configuring an X6 Compute Node

- Provides access to additional firmware and device drivers that might not be applied at installation time, such as adapter cards that are added to the system later

The ServerGuide deploys the OS image to the first device in the boot order sequence. Best practices dictate that you have one device that is available for the ServerGuide installation process. If you boot from SAN, make sure that you have only one path to the installation device because ServerGuide provides no multipath support. See 6.2.6, “Booting from SAN” on page 142 for details.

After the ServerGuide installation procedure, you can attach external storage or activate additional paths to the disk. For installation instructions about how to attach external storage or multipath drivers, see the respective user guide.

The following procedure describes how to install Windows Server 2008 Foundation with ServerGuide. The method to install Linux is similar. Follow these steps:

1. Download the latest version of Lenovo ServerGuide:

   http://ibm.com/support/entry/portal/docdisplay?lndocid=SERV-GUIDE
2. Mount the ServerGuide ISO image through the IMM.
3. Boot the server from the ServerGuide.
4. After the files load, the Start window opens. Choose a language to continue.
5. Select your preferred keyboard layout, and click **Next**.
6. Accept the license agreement, and click **Next**.
7. The Welcome window provides information about the ServerGuide, which systems are supported, a readme file, and copyright and trademark information. Click **Next**.
8. The next window allows you to import additional drivers that are not included in the ServerGuide, as seen in Figure 6-30.

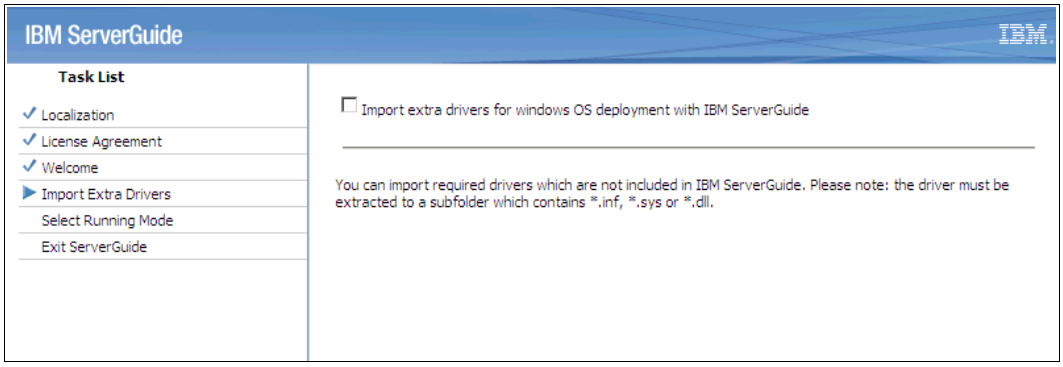

Figure 6-30 Import additional drivers

9. Next, select the operating system that you want to install and click **Next**, as shown in Figure 6-31.

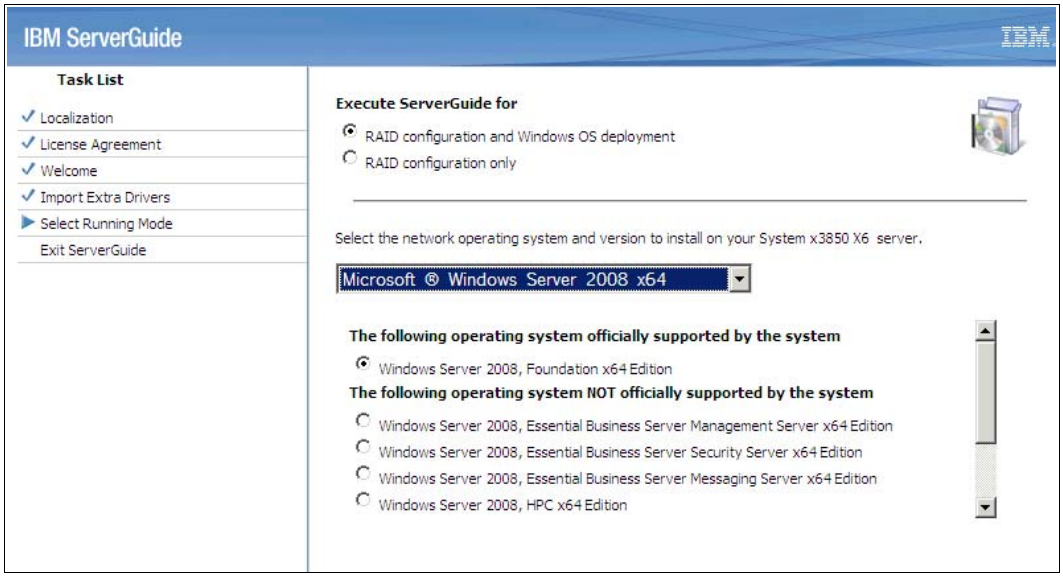

Figure 6-31 Selecting the operating system

10. Enter the current date and time, as shown in Figure 6-32, and click **Next**.

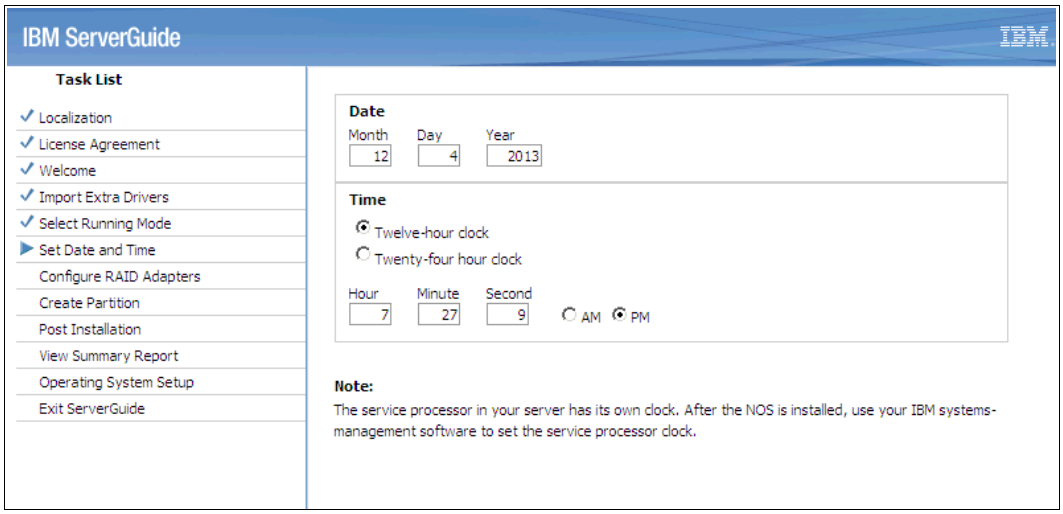

Figure 6-32 Date and time settings

11. Create a RAID configuration. Select a RAID configuration and click **Next**, as shown in Figure 6-33.

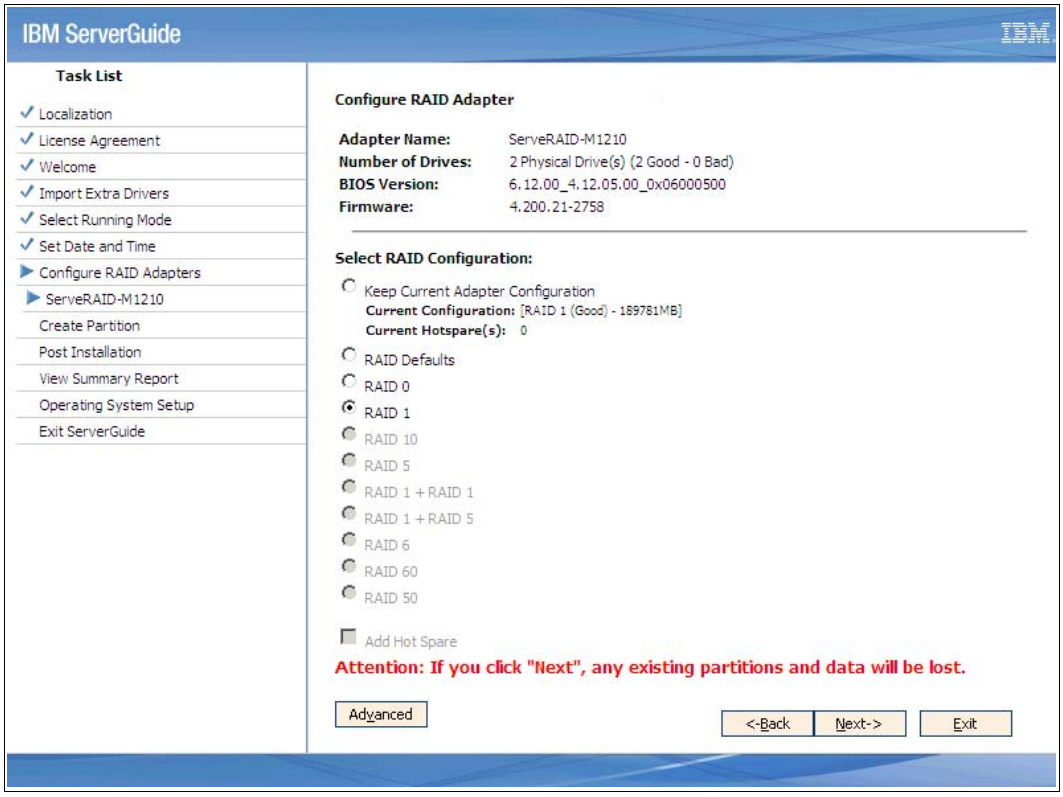

Figure 6-33 RAID selection

12. A confirmation window opens, indicating that the RAID configuration is complete (Figure 6-34). Click **Next**.

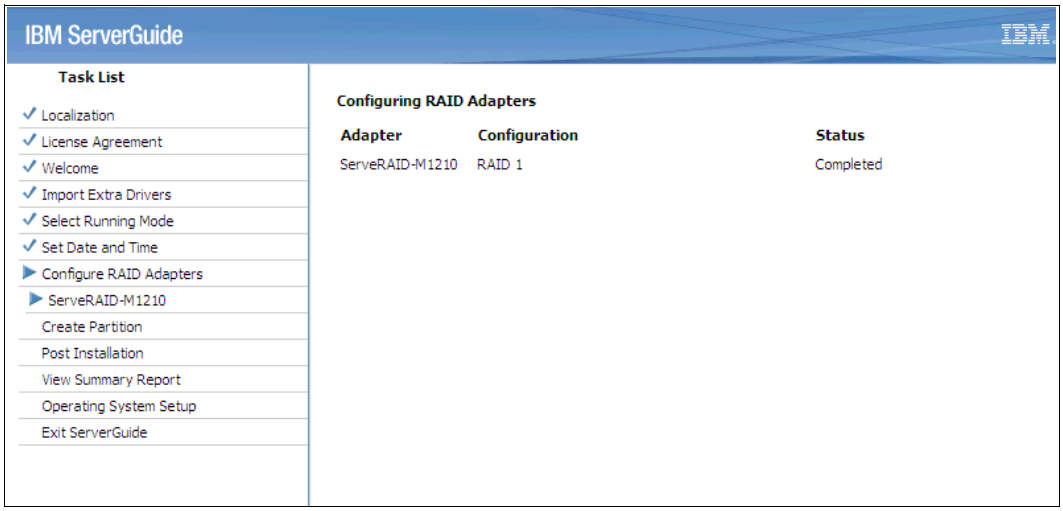

Figure 6-34 RAID confirmation panel

13. You must now create and format a partition. Choose your selection and click **Next** to start the process, as shown in Figure 6-35.

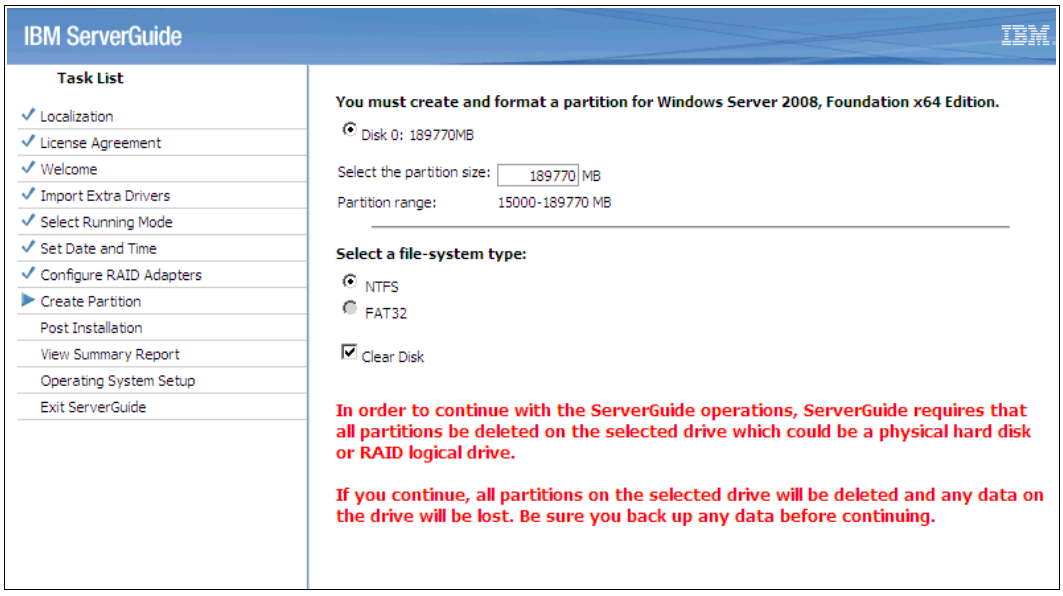

Figure 6-35 Selection for format and partition

14. When the process is finished, click **Next**. You can select postinstallation options, as shown in Figure 6-36.

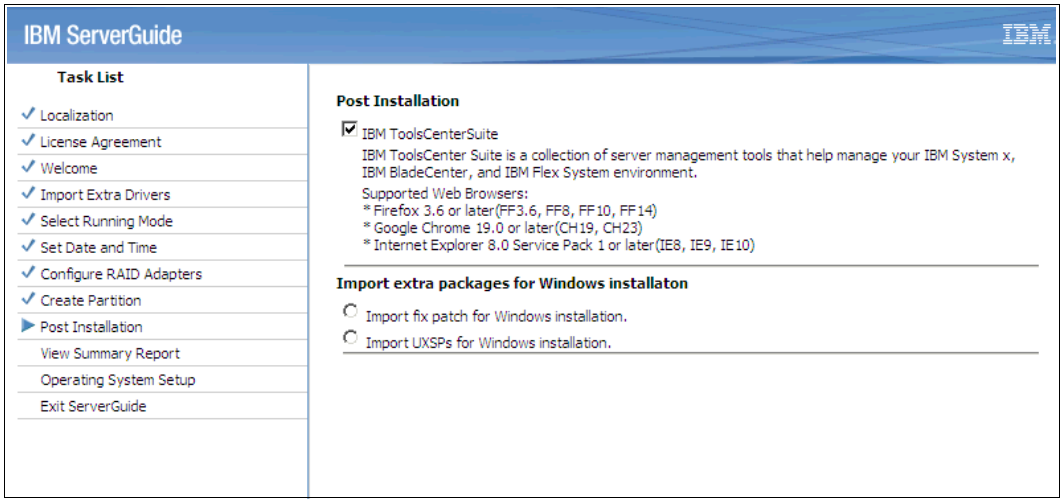

Figure 6-36 Postinstallation options

15. Review the configuration summary, as shown in Figure 6-37. Then, click **Next**.

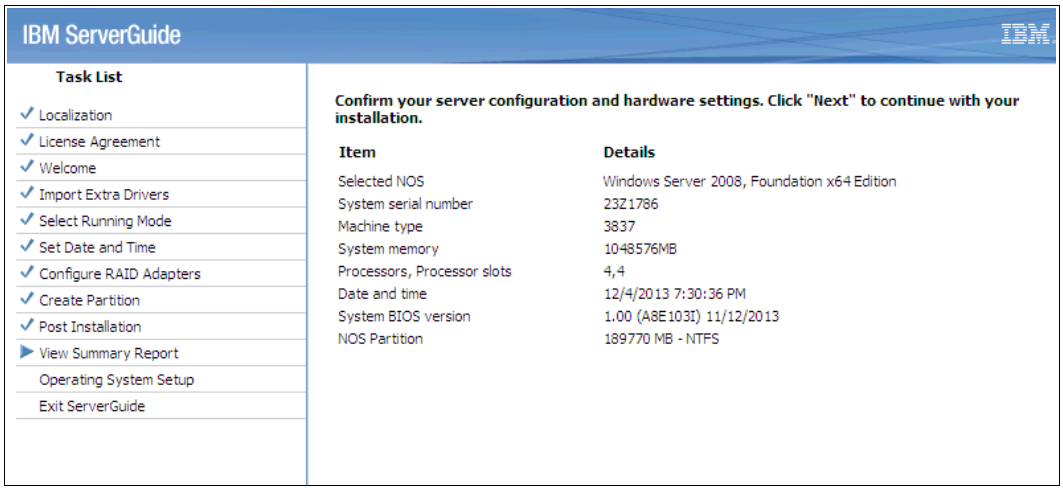

Figure 6-37 Summary report

16. The ServerGuide copies the necessary files to the disk in preparation for the operating system installation, as shown in Figure 6-38.

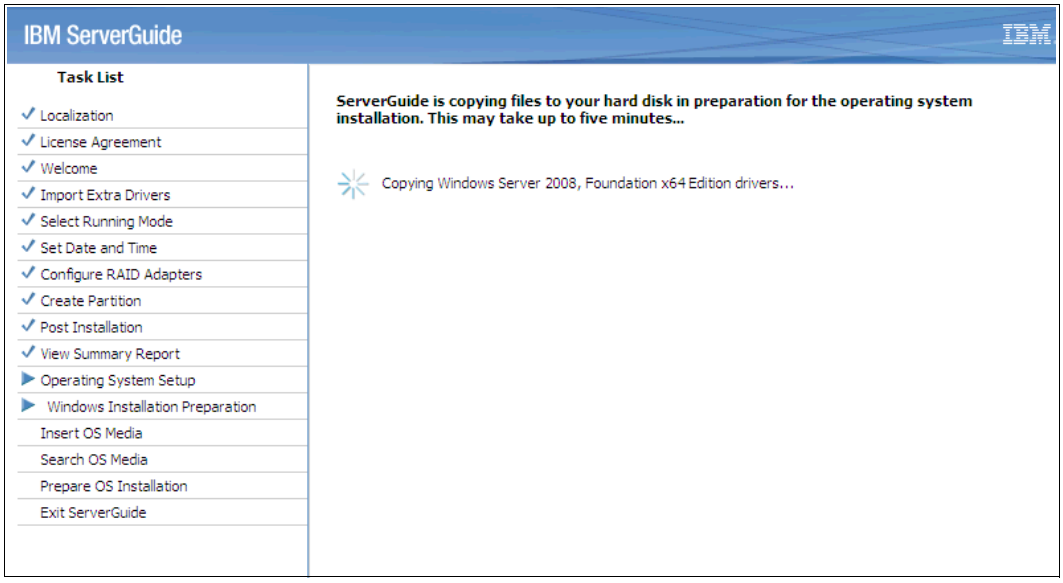

Figure 6-38 ServerGuide copying files

17. When the process is finished, click **Next**.

18. Mount your Microsoft Windows ISO image by using the IMM Virtual Media facility, and then click **Next**, as shown in Figure 6-39. The ServerGuide then searches for the disc.

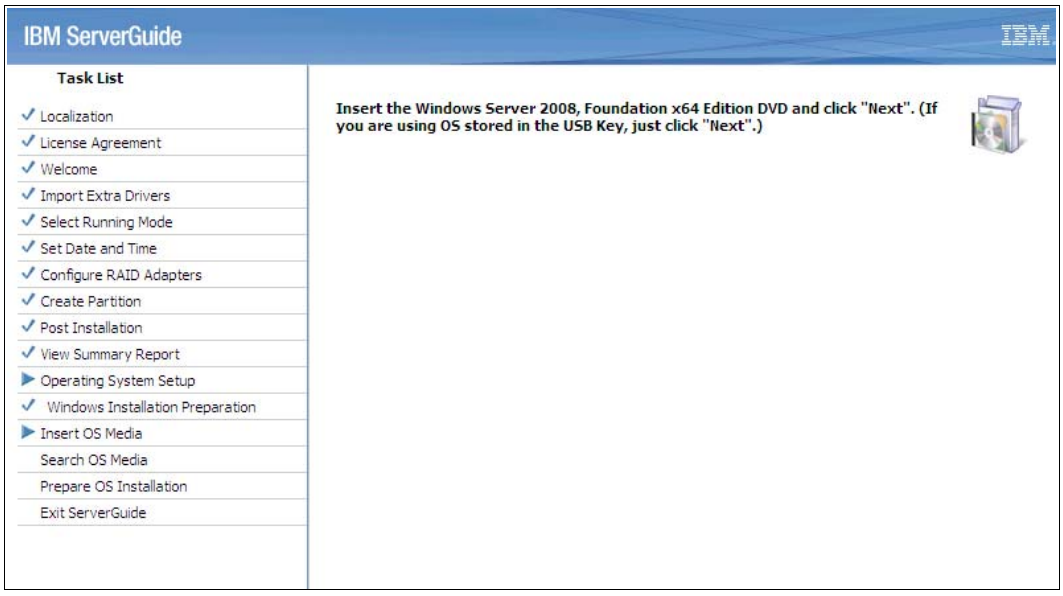

Figure 6-39 Prompt to insert the OS media

19. The Windows setup installation procedure starts. Follow the Microsoft installation procedure to complete the installation of your OS.

6.2.4 Lenovo ServerGuide Scripting Toolkit

You can use the Lenovo ServerGuide Scripting Toolkit to create deployable images by using a collection of system-configuration tools and installation scripts. There are versions of the ServerGuide Scripting Toolkit for the Windows Preinstallation Environment and Linux platforms.

The ServerGuide Scripting Toolkit enables you to tailor and build custom hardware deployment solutions. It provides hardware configuration utilities and OS installation examples for Lenovo Flex System x86-based hardware. If you use it with Lenovo ServerGuide and Lenovo UpdateXpress, the toolkit provides a total solution for deploying Lenovo Flex System x86-based hardware in an unattended mode.

The Scripting Toolkit enables you to create a bootable CD, DVD, or USB key that supports the following tasks and components:

- Network and mass storage devices
- Policy-based RAID configuration
- Configuration of system settings, using the Advanced Settings Utility (ASU)
- Configuration of fibre host bus adapters (HBAs)
- Local self-contained DVD deployment scenarios
- Local CD/DVD and network share-based deployment scenarios
- UpdateXpress System Packs installation integrated with scripted network operating system (NOS) deployment

- IBM Systems Director Agent installation, integration with scripted NOS deployment

The ServerGuide Scripting Toolkit, Windows Edition, enables automated operating system support for the following Windows operating systems:

- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2008, Standard, Enterprise, Datacenter and Web Editions
- Windows Server 2008 x64, Standard, Enterprise, Datacenter and Web Editions
- Windows Server 2008, Standard, Enterprise and Datacenter Editions without Hyper-V
- Windows Server 2008 x64, Standard, Enterprise and Datacenter Editions without Hyper-V
- Windows Server 2008 R2 x64, Standard, Enterprise, Datacenter and Web Editions

The Linux version of the ServerGuide Scripting Toolkit enables automated operating system support for the following operating systems:

- SuSE Linux Enterprise Server 9 32-bit SP4
- SuSE Linux Enterprise Server 9 x64 SP4
- SuSE Linux Enterprise Server 10 32-bit SP1-SP4
- SuSE Linux Enterprise Server 10 x64 SP1-SP4
- SuSE Linux Enterprise Server 11 32-bit Base-SP3
- SuSE Linux Enterprise Server 11 x64 Base-SP3
- Red Hat Enterprise Linux 4 AS/ES 32-bit U6-U8
- Red Hat Enterprise Linux 4 AS/ES x64 U6-U8
- Red Hat Enterprise Linux 5 32-bit U1-U10
- Red Hat Enterprise Linux 5 x64 U1-U10
- Red Hat Enterprise Linux 6 32-bit U1-U5
- Red Hat Enterprise Linux 6 x64 U1-U5

You can download the ServerGuide Scripting Toolkit, a comprehensive list of supported operating systems, or the *Lenovo ServerGuide Scripting Toolkit User's Reference* from the Lenovo ServerGuide Scripting Toolkit web page:

http://ibm.com/support/entry/portal/docdisplay?lndocid=serv-toolkit

6.2.5 Use of embedded VMware ESXi

The X6 Compute Nodes supports a USB flash drive option that is preinstalled with VMware ESXi. VMware ESXi is fully contained on the flash drive, so it does not require any disk space. The Lenovo USB Memory Key for VMware Hypervisor plugs into the internal USB port located on the system board of the X6 Compute Node.

Supported options are listed in 4.14, "Integrated virtualization" on page 67.

VMware ESXi supports booting from the Unified Extensible Firmware Interface (UEFI). To ensure that you can boot ESXi successfully, you must ensure the boot order is correct. The first boot entry must be Embedded Hypervisor. Follow these steps:

1. Press F1 to enter UEFI Setup at POST.
2. Select **Boot Manager** → **Add Boot Option**.
3. Select **Generic Boot Option**, as shown in Figure 6-40 on page 142.

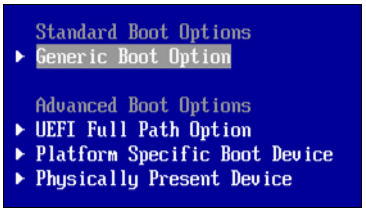

Figure 6-40 Boot options: Generic Boot Option selected

4. Select **Embedded Hypervisor,** as shown in Figure 6-41. If either option is not listed, the option is already in the boot list. When you have finished, press the Esc key to go back to the Boot Manager panel.

Figure 6-41 Boot options: Embedded Hypervisor selected

5. Select **Change the Order** from the Boot Manager menu.
6. Change the boot order to have Embedded Hypervisor at the top, as shown in Figure 6-42.

Figure 6-42 Changing boot order

7. Select **Commit Changes** and press Enter to save changes.

6.2.6 Booting from SAN

Boot from SAN (or *SAN Boot*) is a technique used when the node in the chassis has no local disk drives. It uses an external storage system LUN to boot the OS. Both the operating system and data are on the SAN. This technique is commonly used to provide higher availability and better use of the systems storage (where the operating system is). Hot spare nodes or “Rip-n-Replace” techniques can also be easily implemented by booting from SAN.

To successfully implement SAN Boot, the following conditions must be met:

- Storage system supports SAN Boot.
- Operating system supports SAN Boot.
- FC HBAs, CNAs, or iSCSI initiators support SAN Boot.

SAN booting guidelines

Here are several useful guidelines for booting from SAN:

- Check whether UEFI recognizes the adapter. Select **UEFI** → **System Settings** → **Adapters and UEFI Drivers**. The Adapters and UEFI Drivers panel displays, as shown in Figure 6-43. You need to see Card - HBA. If not, reflash the UEFI, IMM, and firmware of the HBA and check again.

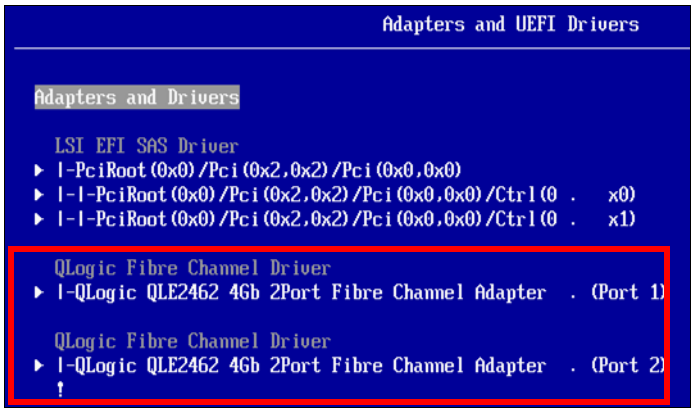

Figure 6-43 Adapters visible in UEFI

- If you do not have internal drives, disable the onboard SAS RAID controller by selecting **System Settings** → **Devices and IO ports** → **Enable/Disable Onboard Devices** and disabling the SAS Controller or Planar SAS.
- Set the HBA as the first device in the Option ROM Execution Order by selecting **System Settings** → **Devices and IO Ports** → **Set Option ROM Execution Order**.
- Only for legacy operating systems that do not support UEFI, set **Legacy Only** as the first boot device.
- Remove all devices from the boot order that do not host an OS. The optimal minimum configuration is CD/DVD and Hard Disk 0. For legacy operating systems only, set **Legacy Only** as the first boot device.
- Enable the BIOS from your HBA.
- Verify that your HBA can see a logical unit number (LUN) from your storage.
- For Microsoft Windows installations, make sure that the LUN is accessible through only one path, either zoning or LUN masking.
- After installation, do not forget to install the multipath driver *before* you set more than one path, if you have more than one path to the LUN.

Sources for more information

You can also check SAN support and requirements and storage vendors for documentation for the operating system that is used to boot from. See the following sources for more SAN boot-related information:

- IBM System Storage Interoperability Center for compatibility information:

  `http://www.ibm.com/systems/support/storage/config/ssic`

- *Red Hat Enterprise Linux 6 Installation Guide*, “Storage Devices” chapter:

  `http://red.ht/1v6Dmai`

- *vSphere Storage* document, available to download from VMware Publications:

  `http://vmw.re/1AF57an`

- *Windows Boot from Fibre Channel SAN – Overview and Detailed Technical Instructions for the System Administrator,* available from the Microsoft Download Center:

  http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=2815

7

Management

The Lenovo Flex System Enterprise Chassis brings a whole new approach to management. This approach is based on a global management appliance, the IBM Flex System Manager, which you can use to view and manage functions for all of your Enterprise Chassis components. These components include the Chassis Management Module (CMM), I/O modules, computer nodes, and storage. The Flex System Manager is standard with PureFlex System configurations and is optional with Lenovo Flex Systems.

Systems management at the basic chassis level uses the CMM and the native switch managers that are supplied by each I/O module.

Management of the Enterprise Chassis with the CMM and Flex System Manager provides the most comprehensive management of the chassis and all components. Other functions, such as VM Control, Storage Management, Update Manager, and operating systems monitoring and management are also included in this combination.

This chapter includes the following topics:

7.1 Management network

The Lenovo Flex System Enterprise Chassis provides separate management and data networks. The management network is a private and secure Gigabit Ethernet network that is used to perform management-related functions throughout the chassis, including management tasks on compute nodes, switches, and the chassis. The data network normally is used for operating system and user access administration and applications.

The management network connection is externalized only through the CMM's network connection. The data network is externalized through the external switch ports of the switch I/O modules. These switches and switch ports can be configured by using traditional methods.

The management network is shown in Figure 7-1 (blue lines). It connects the CMM to the compute nodes' Integrated Management Module (IMM), the management connection on the switches in the I/O bays, and the Flex System Manager eth0 management interface. The Flex System Manager connection to the management network is through a special Broadcom 5718-based management network adapter. The management networks in multiple chassis are connected through the external ports of the CMMs in each chassis through a GbE top-of-rack switch.

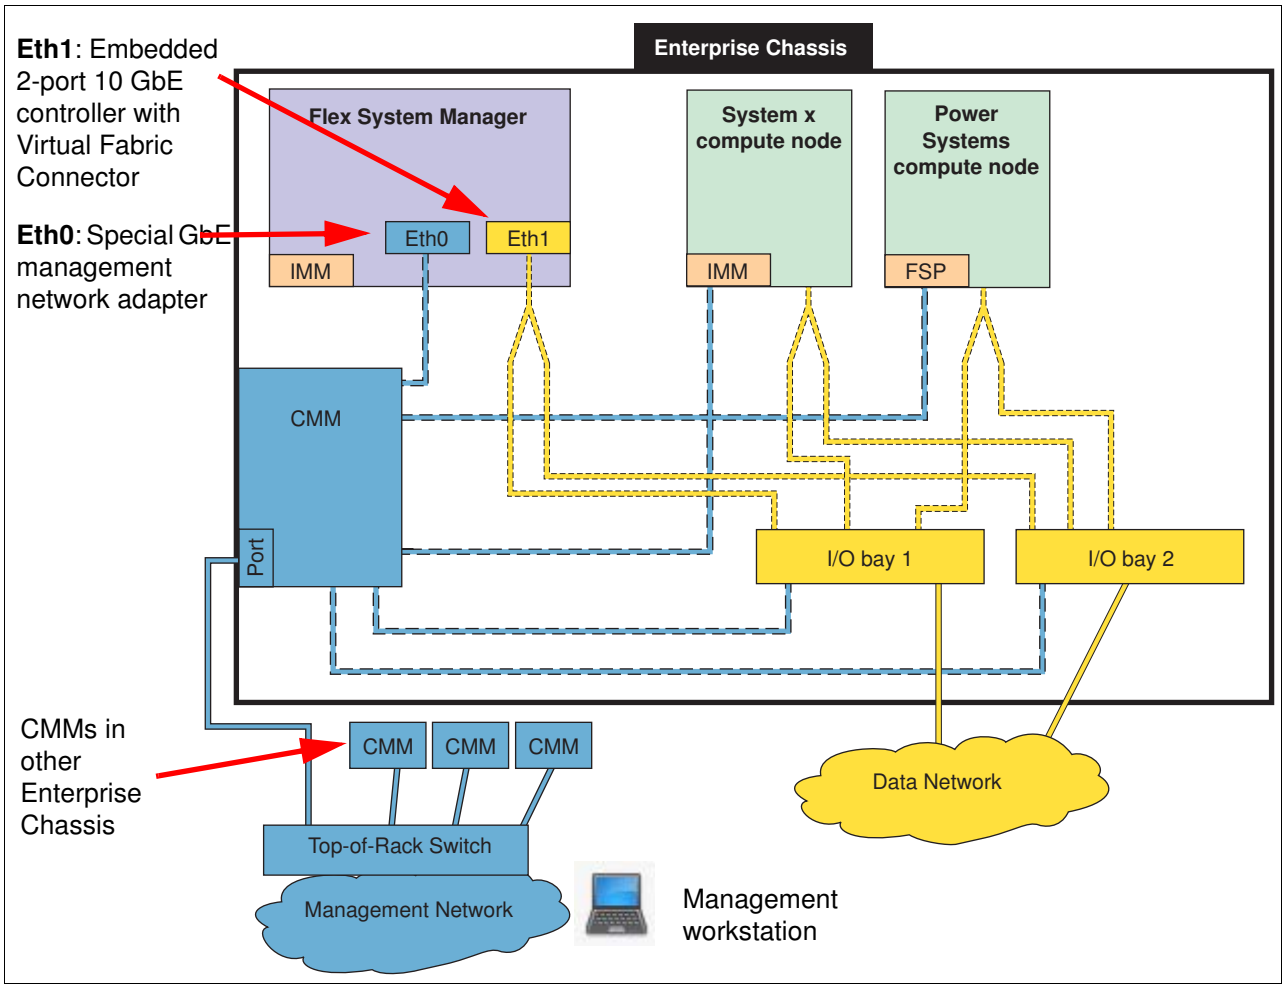

Figure 7-1 Separate management and production data networks

The yellow lines in Figure 7-1 on page 146 show the production data network. The Flex System Manager also connects to the production network through its eth1 interface on the adapter. This can be set up to access the Internet (proxy or direct) for product firmware and software updates and for dial-home functionality.

PureFlex System and IPv6: In a PureFlex System configuration, all components on the management network are configured with a static IPv6 address with the prefix of `fd8c:215d:1783:c0de`, including eth0 on the Flex System Manager. In addition, the eth0 Flex System Manager interface does not get an IPv4 address. Normal access to the Flex System Manager user interface is through an IPv4 address that is assigned to eth1.

One of the key functions that the data network supports is discovery of operating systems on the various network endpoints. This discovery by the Flex System Manager is required to support software updates and compliance on an endpoint, such as a compute node. You can use the Flex System Manager Checking and Updating Compute Nodes wizard to discover operating systems as part of the initial setup.

7.2 Chassis Management Module (CMM)

This section gives a brief overview of the CMM, as shown in Figure 7-2.

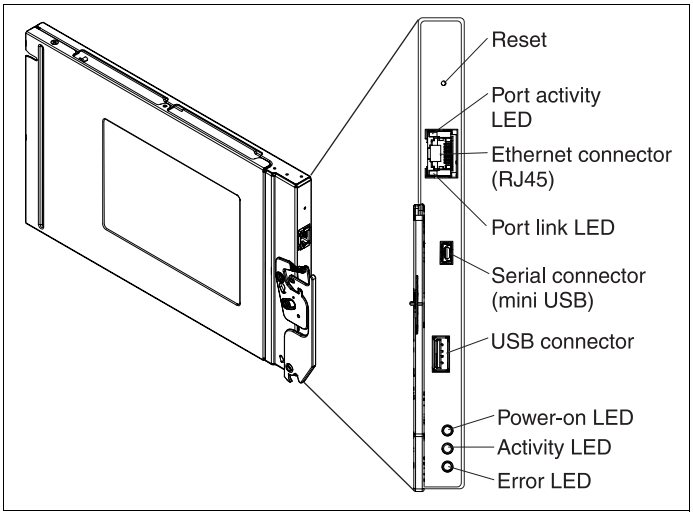

Figure 7-2 Chassis Management Module

Detailed CMM setup and overall use information is not covered in this book. For more information, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060:

`http://lenovopress.com/sg248060`

For a hardware overview of the CMM, see *Flex System Products and Technology*, SG24-8255:

`http://lenovopress.com/sg248255`

7.2.1 CMM overview

The CMM is a hot-swap module that provides single-chassis management and is used to communicate with the management controller in each compute node. It provides system

monitoring, event recording, and alerts, and it manages the chassis, its devices, and the compute nodes. The chassis supports up to two CMMs. If one CMM fails, the second CMM (if present) can detect its inactivity, self-activate, and take control of the system without any disruption. The CMM is central to the management of the hardware in the chassis.

The CMMs are inserted in the back of the chassis and are vertically oriented. When you are looking at the back of the chassis, the CMM bays are on the far right side, as shown in Figure 7-3. CMM bay 1 is the lower position, and CMM 2 is the upper position.

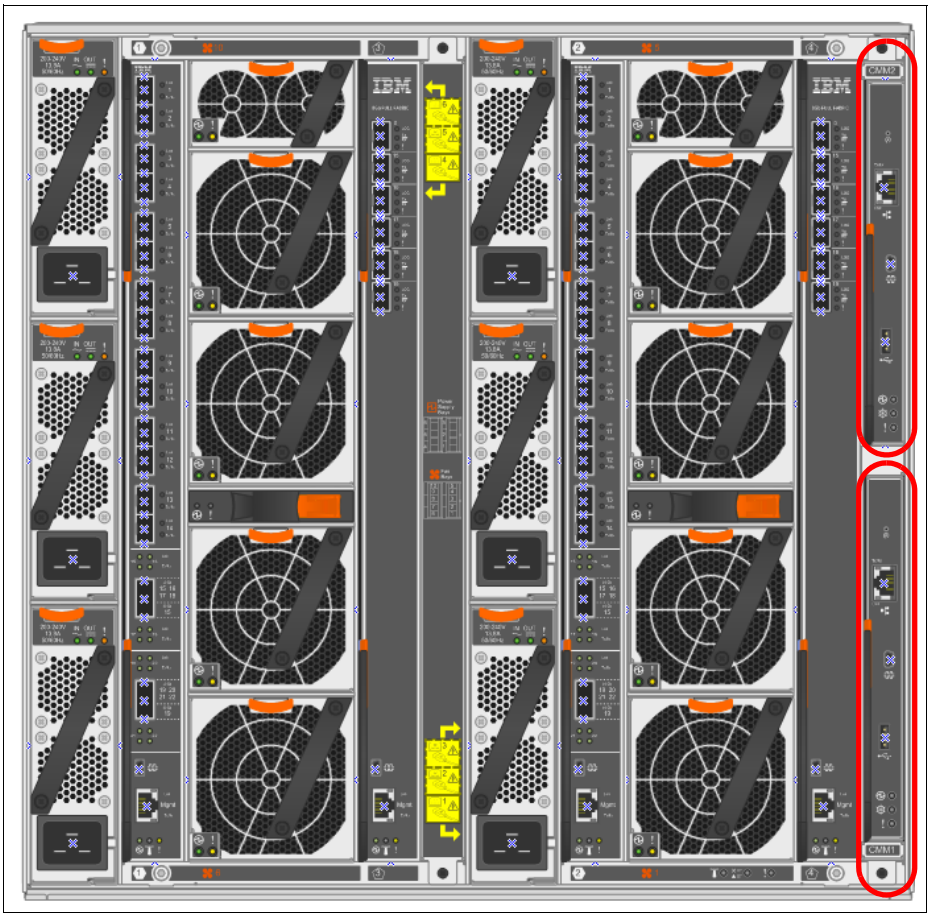

Figure 7-3 Chassis Management Module bays

Through an embedded firmware stack, the CMM implements functions to monitor, control, and provide external user interfaces to manage all chassis resources. You can use the CMM to perform the following functions:

- Define login IDs and passwords
- Configure security settings, such as data encryption and user account security
- Select recipients for alert notification of specific events
- Monitor the status of the compute nodes and other components
- Find chassis component information
- Discover other chassis in the network and enable access to them
- Control the chassis, compute nodes, and other components
- Access the I/O modules to configure them
- Change the start sequence in a compute node
- Set the date and time
- Use a remote console for the compute nodes
- Enable multi-chassis monitoring

- Set power policies and view power consumption history
- Access support for Lenovo Features on Demand (FoD)
- Access support for Lenovo Fabric Manager

The CMM automatically detects installed switches, compute nodes, and storage nodes in the Enterprise Chassis and stores vital product data on them.

7.2.2 CMM user interfaces

The CMM supports a web-based graphical user interface (GUI) that provides a way to perform chassis management functions within a supported web browser. You can also perform management functions through the CMM command-line interface (CLI). Both the GUI and CLI are accessible through the single RJ45 Ethernet connector on the CMM or from any system that is connected to the same network.

The default security setting is `Secure`, so HTTPS or SSH is required to connect to the CMM.

7.2.3 CMM default network information

By default, the CMM is configured to respond to the Dynamic Host Configuration Protocol (DHCP) first before a static IPv4 address is used as a fallback. If a DHCP response is not received within 3 minutes of the CMM Ethernet port connecting to the network, the CMM uses the factory default IP address and subnet mask. During this 3-minute interval, the CMM is inaccessible. The IP behavior can be changed during the initial setup with a locally attached workstation.

A new CMM or a CMM that is reset through the “pinhole” has the following default settings:

- IP address: DHCP; if no response, then 192.168.70.100
- Subnet mask: 255.255.255.0
- User ID: USERID (case sensitive)
- Password: PASSW0RD (case sensitive, with a zero rather than the letter O and requires changing on first use)

PureFlex System defaults: For PureFlex System configurations, the following default settings are used:

- Static IP address (DHCP off)
- IP address:
  - 192.168.93.100 (first chassis)
  - 192.168.93.102 (second chassis)
  - 192.168.93.104 (third chassis)
- Subnet mask: 255.255.252.0
- User ID: USERID (case sensitive)
- Password: PASSW0RD (case sensitive, with a zero rather than the letter O, and requires changing upon first use)

A “pinhole” reset of the CMM in a PureFlex configuration reverts the CMM to the non-PureFlex defaults.

7.2.4 CMM requirements

At least one CMM is required for each chassis for control and management (a second CMM is optional but recommended for resiliency reasons).

The CMM and compute node management modules (the IMM), storage nodes (IMM), or I/O modules are required to be on the same subnet.

7.3 IBM Flex System Manager

This section gives a brief overview of Flex System Manager, as shown in Figure 7-4.

Figure 7-4 Flex System Manager

Detailed Flex System Manager setup and overall use information is not covered in this book, but that information is available in the Lenovo Press publication *Implementing Systems Management of IBM PureFlex System*, SG24-8060:

http://lenovopress.com/sg248060

7.3.1 Flex System Manager overview

The Flex System Manager is a high-performance, scalable system management appliance that is based on the Lenovo Flex System x240 Compute Node. Flex System Manager hardware has systems management software preinstalled, and you can configure, monitor, and manage Flex System Manager resources in up to 16 chassis.

The Flex System Manager looks similar to the x240 Compute Node internally. However, there are differences that make these two hardware nodes not interchangeable.

From a hardware perspective, Flex System Manager is a locked-down compute node with a single specific hardware configuration that is designed for optimal performance of the preinstalled software stack. This hardware configuration currently includes an eight-core 2.0 GHz processor, 32 GB of RAM, two 200 GB sold-state drives (SSDs) in a RAID-1 configuration, and a 1 TB hard disk drive (HDD).

A management network adapter is a standard feature of the Flex System Manager and provides a physical connection into the private management network of the chassis. This card is one of the features that makes the Flex System Manager unique when it is compared to other Flex System Compute Nodes. The management network adapter provides a physical

connection to the private management network of the chassis so that the software stack has visibility into the data and management networks.

The preinstallation contains a set of software components that are responsible for performing certain management functions. These components must be activated by using the available FoD software entitlement licenses, and they are licensed on a per-chassis basis. You need one license for each chassis that you plan to manage.

The management node comes standard without any entitlement licenses, so you must purchase a license to enable the required Flex System Manager functionality. There are two versions of IBM Flex System Manager: base and advanced.

PureFlex note: In a PureFlex configuration, *Flex System Manager base* is included as part of the configuration and is licensed for the total number of chassis that are included in the original order. *Flex System Manager advanced* is optional in all PureFlex configurations.

The Flex System Manager base feature set offers the following functionality:

- Support for up to 16 managed chassis
- Support for up to 8000 managed endpoints
- Auto-discovery of managed elements
- Overall health status reports
- Monitoring and availability
- Hardware management
- Security management
- Administration
- Network management (Network Control)
- Storage management (Storage Control)
- Virtual machine lifecycle management (VMControl Express)

The Flex System Manager Advanced feature set offers all of the capabilities of the base feature set plus the following features:

- Image management (VMControl Standard)
- Pool management (VMControl Enterprise)

Flex System Manager management software includes the following features:

- Monitoring and problem determination:
  - A real-time multi-chassis view of hardware components with overlays for more information
  - Automatic detection of issues in your environment through event setup that triggers alerts and actions
  - Identification of changes that might affect availability
  - Server resource use by a virtual machine or across a rack of systems
- Hardware management:
  - Automated discovery of physical and virtual servers and interconnections, applications, and supported third-party networking.
  - Inventory of hardware components
  - Chassis and hardware component views
  - Hardware properties
  - Component names and hardware identification numbers

- Firmware levels
- Usage rates

- Network management:
  - Management of network switches from various vendors
  - Discovery, inventory, and status monitoring of switches
  - Graphical network topology views
  - Support for Keyboard, Video, and Mouse (KVM) switches, VMware virtual switches, and physical switches
  - VLAN configuration of switches
  - Integration with server management
  - Per-virtual machine network use and performance statistics that are provided to VMControl
  - Logical views of servers and network devices that are grouped by subnet and VLAN
- Storage management:
  - Discovery of physical and virtual storage devices
  - support for virtual images on local storage across multiple chassis
  - Inventory of physical storage configuration
  - Health status and alerts
  - Storage pool configuration
  - Disk sparing and redundancy management
  - Virtual volume management
  - Support for virtual volume discovery, inventory, creation, modification and deletion
- Virtualization management (base feature set)
  - Support for VMware, Hyper-V, KVM and IBM PowerVM
  - Create virtual servers
  - Edit virtual servers
  - Manage virtual servers
  - Relocate virtual servers
  - Discover virtual server, storage, and networking resources, and virtualize the physical-to-virtual relationships
- Virtualization management (advanced feature set):
  - Create image repositories for storing virtual appliances and discover existing image repositories in your environment
  - Import external, standard-based virtual appliance packages into your image repositories as virtual appliances
  - Capture a running virtual server that is configured the way that you want, complete with guest operating system, running applications, and virtual server definition
  - Import virtual appliance packages that exist in the Open Virtualization Format (OVF) from the Internet or other external sources
  - Deploy virtual appliances quickly to create virtual servers to meet business need
  - Create, capture, and manage workloads

- Create server system pools, where you can consolidate your resources and workloads into distinct manageable groups
- Deploy virtual appliances into server system pools
- Manage server system pools, including adding hosts or more storage space and monitoring the health of the resources and the status of the workloads in them
- Manage and group storage subsystems into pools

- Additional features:
  - A resource-orientated chassis map provides an instant graphical view of chassis resources, including nodes and I/O modules:
    - Provides an instant view of individual system status and inventory
    - A chassis map provides an inventory view of active statuses that require administrative attention and a compliance view of system firmware
    - Actions can be taken on nodes, such as working with server-related resources, showing and installing updates, submitting service requests, and starting the remote access tools
  - Remote console:
    - Remote control session from any connected web browser
    - Remove media support to allow the use of a local optical drive or a local ISO file
    - Remote power operations (power on, off, and restart)
  - Hardware detection and inventory creation
  - Firmware compliance and updates
  - Automation detection of hardware failures:
    - Provides alerts
    - Takes corrective action
    - Notifies Lenovo of problems to escalate problem determination (if the IBM Electronic Service Agent is enabled)
  - Reports of health status (such as processor use) on all hardware devices from a single chassis view
  - Administrative capabilities, such as setting up users within profile groups, assigning security levels, security governance

7.3.2 Flex System Manager user interfaces

The Flex System Manager supports a web-based GUI that provides access to all Flex System Manager management functions from a supported web browser. You can also perform management functions through the Flex System Manager CLI. The GUI and CLI should be available through a network connection after the Flex System Manager setup wizard completes.

The default security setting is `Secure`, so HTTPS or SSH is required to connect to the Flex System Manager.

7.3.3 Flex System Manager requirements

The Flex System Manager requires one compute node bay slot in the chassis. When the Flex System Manager is installed into an empty slot, all connections to the chassis management

and data networks are made automatically through the midplane of the chassis to the CMM and I/O switches.

After the Flex System Manager is installed in the chassis and discovered by the CMM, the Flex System Manager setup wizard must be run. The setup wizard requires a virtual console through the Flex System Manager IMM remote console facility or through a KVM connected through the KVM breakout dongle connected to the front of the Flex System Manager. The Flex System Manager setup wizard starts automatically during the initial boot of the Flex System Manager.

For more information about this process, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060, which is available from:

http://lenovopress.com/sg248060

7.4 Comparing CMM and Flex System Manager management

The CMM is a mandatory requirement within the Enterprise Chassis, covering physical infrastructure management. In addition to interacting with the CMM to perform infrastructure management, the Flex System Manager provides management of components and functions, such as virtualization hypervisors, operating systems, software patch management, and compliance.

Table 7-1 on page 155 compares the management functions that are available in the CMM and the Flex System Manager (abbreviated as FSM in the UI, paths, and files).

Table 7-1 CMM and Flex System Manager management capabilities comparison

Capability	CMM	FSM
Web-based GUI	Yes	Yes
CLI	Yes	Yes
Number of chassis managed	1	16
Number of compute nodes managed	14	225
Activate or shut down virtual machines	No	Yes
NPIV (N_Port ID Virtualization) management	No	Yes
Suspend or resume	No	Yes
Shared storage pools	No	Yes
Shared network pools	No	Yes
Shared compute pools	No	Yes
System firmware updates	No	Yes
Concurrent system firmware updates	No	Yes
Adapter updates	No	Yes
Operating system updates	No	Yes
Cloud-enabled	No	Yes
Light path information	Yes	Yes
Monitors and Event action plans	No	Yes
Service focal point and Call home	Yes	Yes

7.5 Management by CMM

This section describes the basic steps of managing an X6 Compute Node from the CMM.

7.5.1 Accessing the CMM

Before you begin, you need the IP address of the CMM. You can access the CMM by using SSH or a web browser. The browser method is described here.

Complete the following steps:

1. Open a web browser and navigate to the following URL (where `systemize` is the host name or IP address of the CMM). The protocol used is *https*, not http:

   `https://systemize`

The login window that is shown in Figure 7-5 opens.

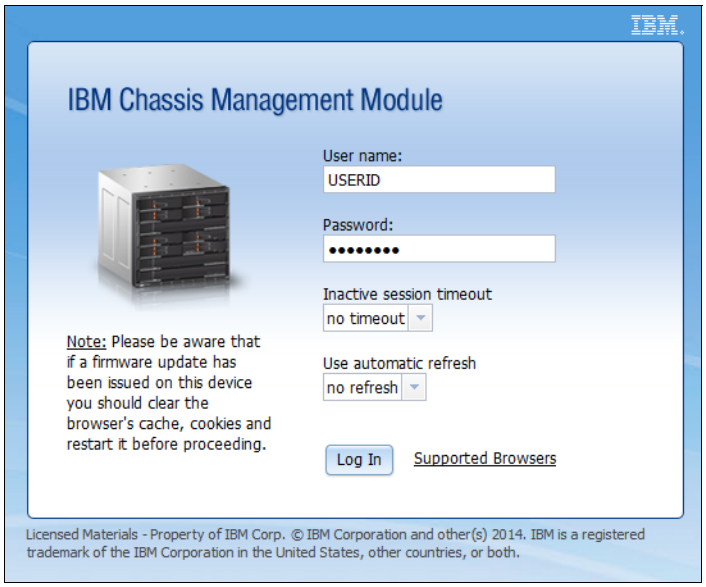

Figure 7-5 CMM login window

2. Log in with your user ID and password. The System Status window of the CMM opens with the Chassis tab active, as shown in Figure 7-6. If you don't see the chassis view, click **System Status** from the menu bar at the top of the window.

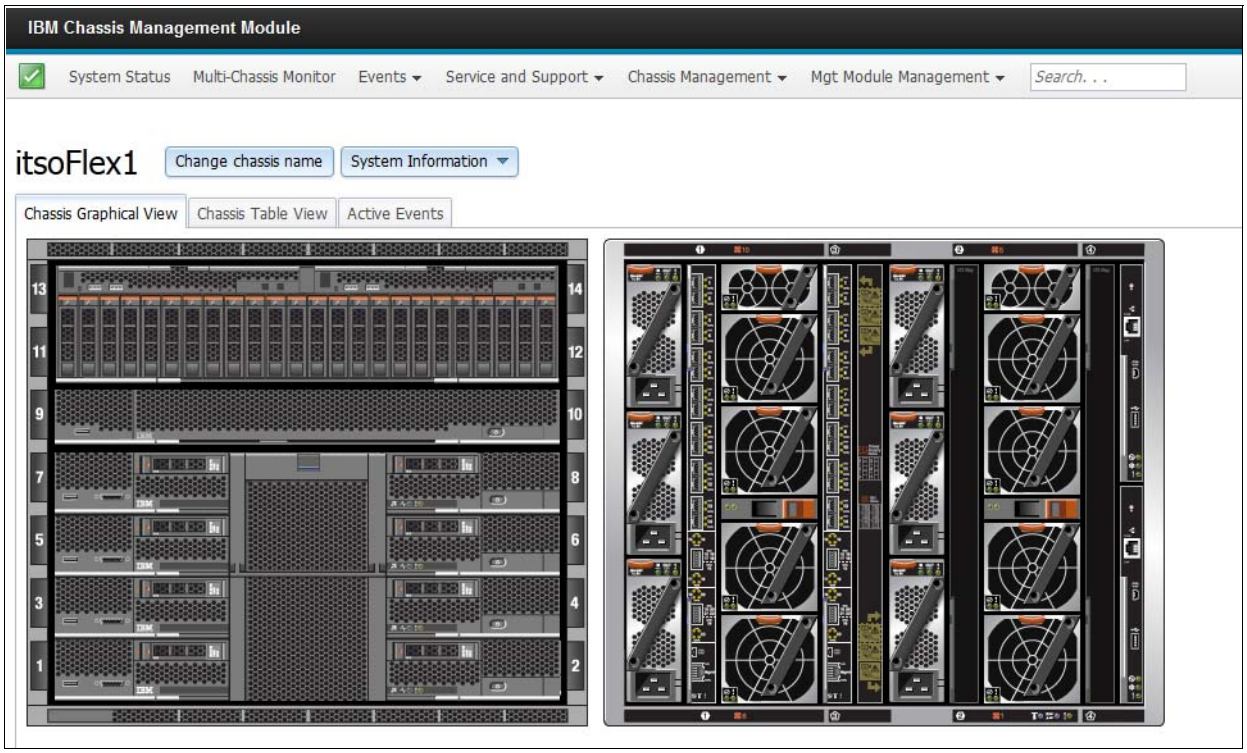

Figure 7-6 CMM initial view: System Status

The CMM web interface has a navigation menu structure at the top of each page that gives you easy access to most functions, as shown in Figure 7-7. Most menu options display more functions when clicked.

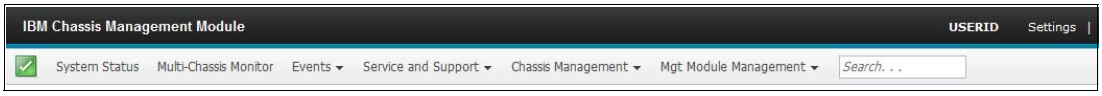

Figure 7-7 CMM top menu bar

The following navigation drop-down menus are available from the bar:

- System Status
- Multi-Chassis Monitor
- Events, as shown in Figure 7-8

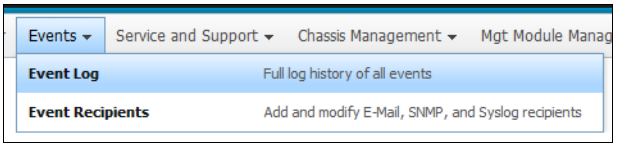

Figure 7-8 CMM Events drop-down menu

- Service and Support, shown in Figure 7-9.

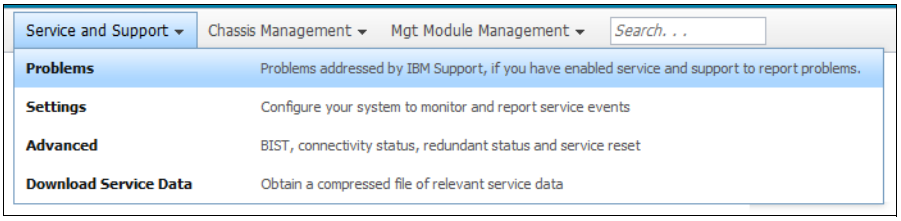

Figure 7-9 CMM Service and Support drop-down menu

- Chassis Management, shown in Figure 7-10.

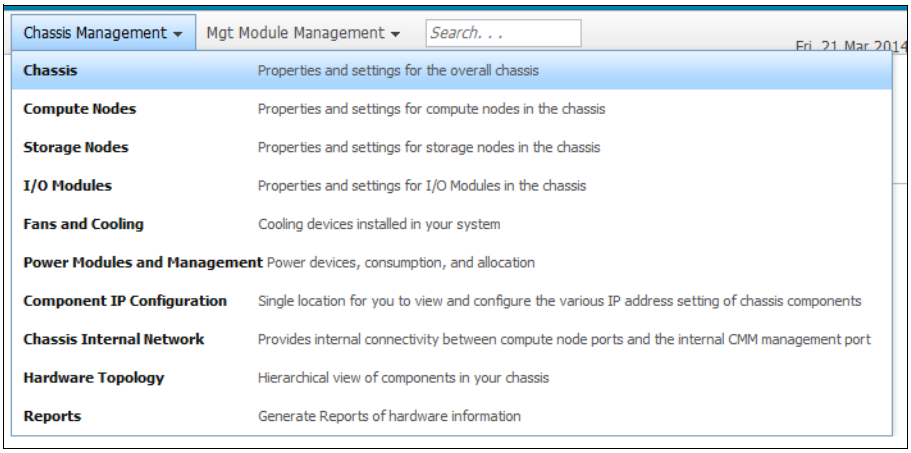

Figure 7-10 CMM Chassis Management drop-down menu

- Management Module Management, shown in Figure 7-11.

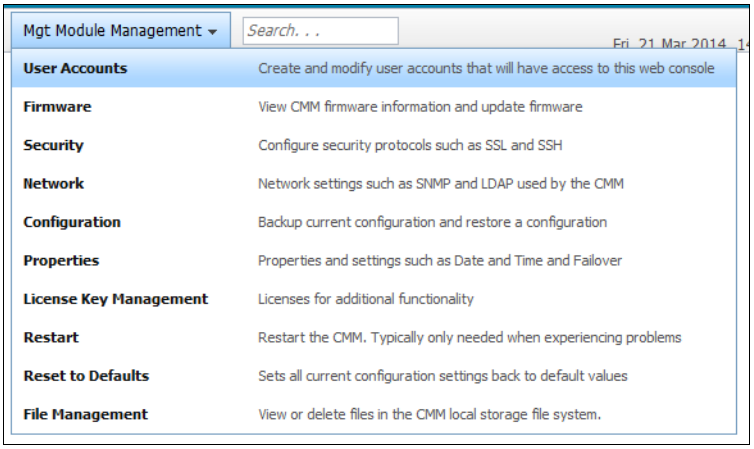

Figure 7-11 Management module management drop-down menu

The following options are of most interest for managing compute nodes and are described in in 7.5.3, “X6 Compute Node management” on page 159.

- System Status
- Chassis Management → Compute Nodes
- Chassis Management → Component IP Configuration

The Service and Support tab information is described in 7.5.4, “Service and Support option” on page 164.

7.5.2 Connecting an X6 Compute Node to the CMM

During a chassis power up or when a compute node is first inserted into the chassis, the CMM automatically performs a discovery process that detects and collects information about the new system. No other action is required to connect or register the new compute node to the CMM.

This process is indicated on a newly inserted compute node by a fast blink of the green power indicator LED. When the discovery process is complete, the LED changes to a slow blink and actions can be performed on the compute node. The discovery process for nodes can take approximately 120 seconds.

During the discovery process, the System Status view (Figure 7-12) gives a visual indication of a node in a discovery status.

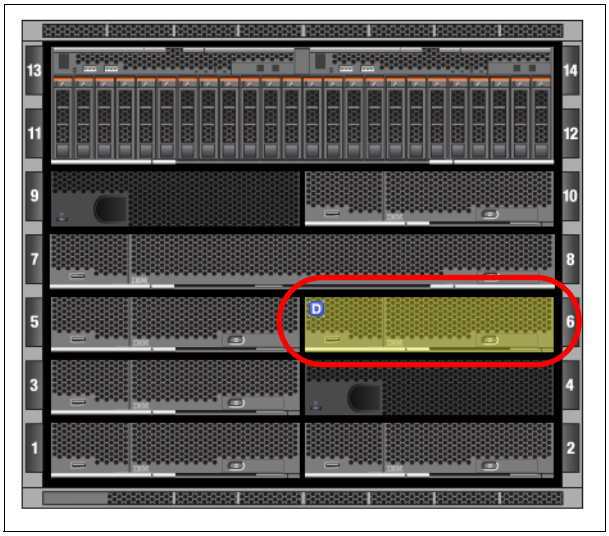

Figure 7-12 Bay 6 compute node in discovery status

7.5.3 X6 Compute Node management

This section describes X6 Compute Node management options through the CMM and how to use these options to allow management by more advanced platform managers.

When you are performing management operations on x86 or on IBM Power based compute nodes, there are two primary places at the top of the page menu structure of the CMM that are used: the *System Status* tab and the *Chassis Management* tab.

System Status option

The System Status option shows a graphical chassis map window, which is the default view when you enter the CMM web interface. You can also access it by clicking **System Status**.

The chassis map is active and shows changes in status of the chassis components by changes in colors and various symbols. Placing the mouse cursor over a component shows vital product data, such as model, type, serial number, and general health status.

The chassis map is also interactive and allows the selection of a component to display the available actions, such as power on/off, boot options, and location LEDs.

Below the actions, a detail window shows all of the available information for a chassis component that is categorized by a row of tabs. These details are read-only from the Chassis View tab, but user-changeable options can be modified by navigating from the top menu: **Chassis Management** → **Compute Node**.

Figure 7-13 shows the active and interactive modes on the System Status view.

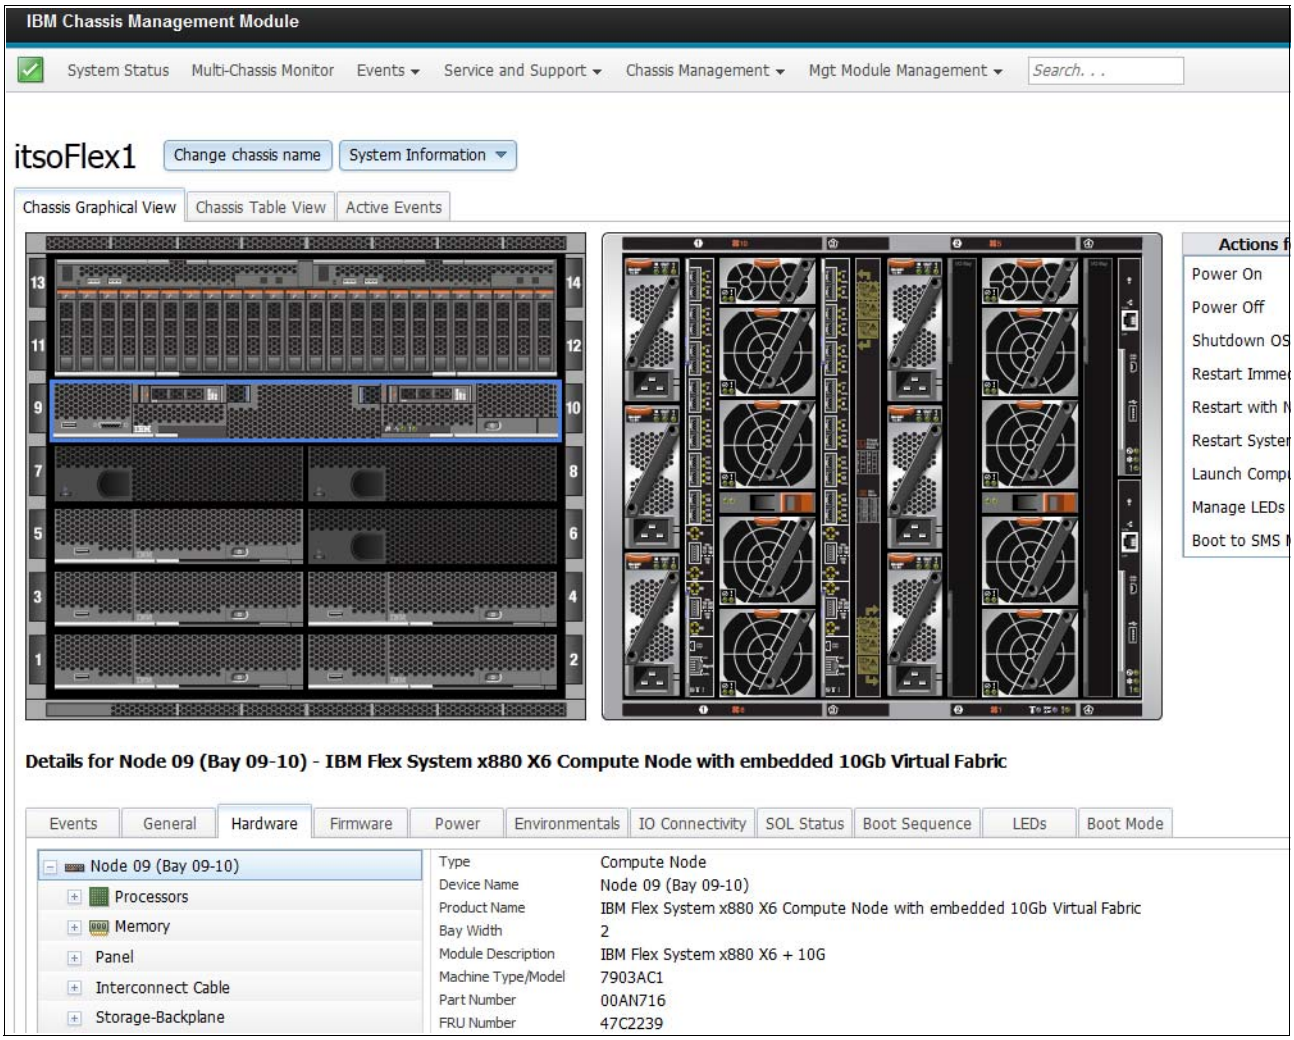

Figure 7-13 x880 compute node management options from System Status view

Component IP Configuration option

The automatic node discovery process of the CMM allows the basic management by the CMM without further configuration. IP address configuration of the individual nodes is required if you want management by the native interface of the node or an advanced manager, such as the Flex System Manager.

The Component IP Configuration option is used to configure the IP addresses for the I/O modules, compute nodes, and storage nodes. These IP addresses are required to be in the same subnet as the CMM. The switch function of the CMM provides the connectivity for each IMM and service processor of the different node types from the chassis management network to an external network. This network traffic flows through the CMM's external 1Gb Ethernet connection.

The ability of the Flex System Manager to manage a X6 Compute Node is dependent on communicating with the IMM. Proper configuration of the IMM IP information is also required to access the IMM's web interface or CLI.

Configuration of these components is started by clicking **Chassis Management** → **Component IP Configuration**, which displays the page, as shown in Figure 7-14.

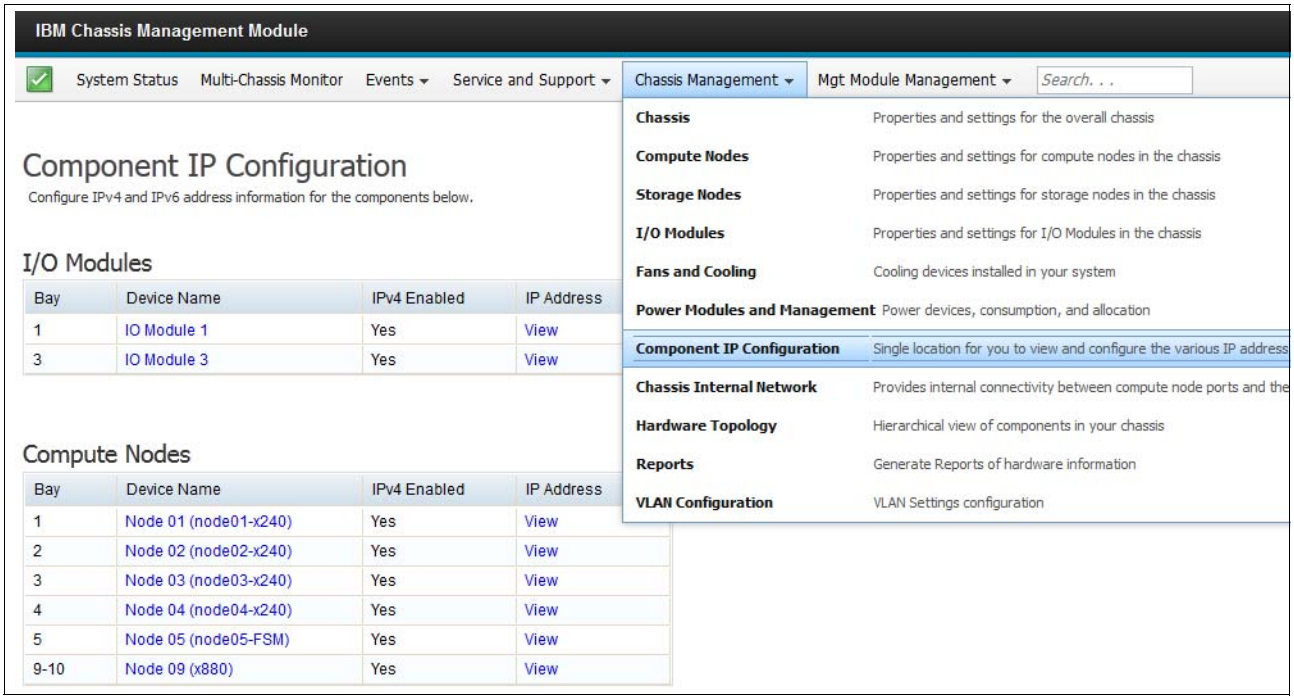

Figure 7-14 Component IP configuration

From this view, the IP configuration information of the I/O modules, compute nodes, and storage nodes can be reviewed by clicking the **View** option for the node that you want informatino about, as shown in Figure 7-15

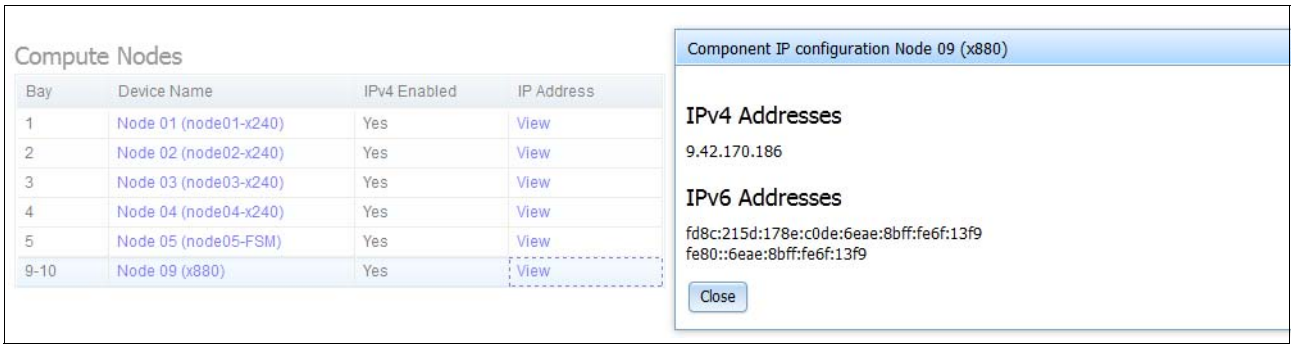

Figure 7-15 Reviewing the current node IP configuration with the View option

To edit or configure the IPv4 and IPv6 addresses, you click the entry in the Device Name column (as shown in Figure 7-15 on page 161) of the node that you want informatin about and then the appropriate tab in the configuration window, as shown in Figure 7-16. Enter the network configuration information that you want and click **Apply**.

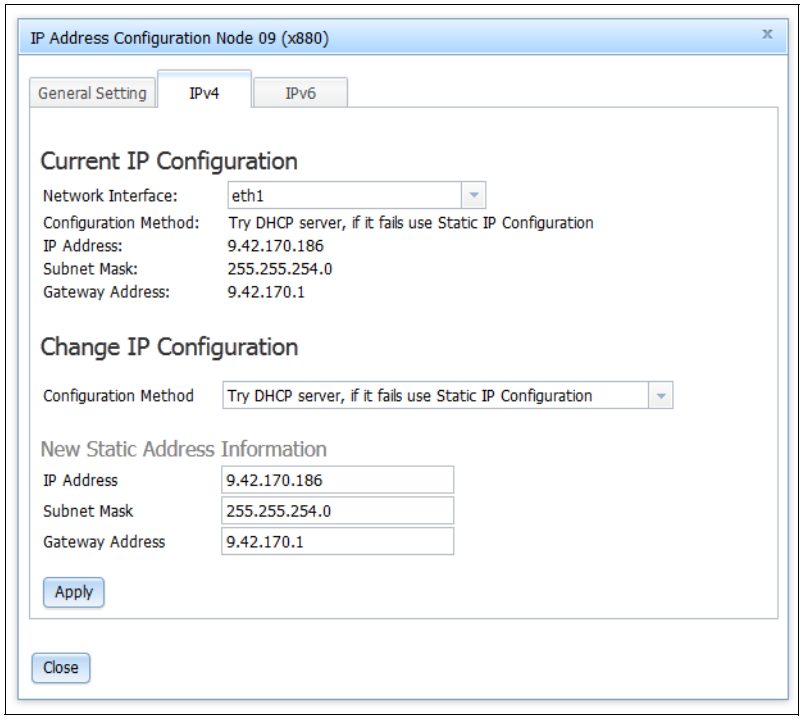

Figure 7-16 Configuring IPv4 information for an IMM

Figure 7-17 shows the confirmation message that settings have been applied. Click **Close** on the confirmation message and **Close** on the configuration window to return to the component IP Configuration page.

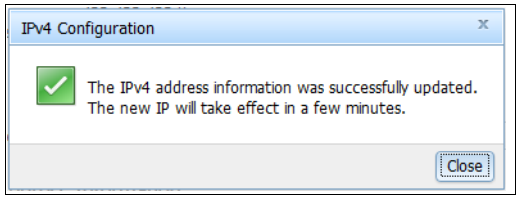

Figure 7-17 IP configuration change confirmation

The configuration changes take several moments to occur, and the component IP Configuration view must be manually refreshed to update the information.

Compute node management

Clicking **Chassis Management** → **Compute Nodes** (as shown in Figure 7-18) displays a list of all compute nodes that are installed in the chassis. The Device Name column contains active links; the remaining columns are for information only.

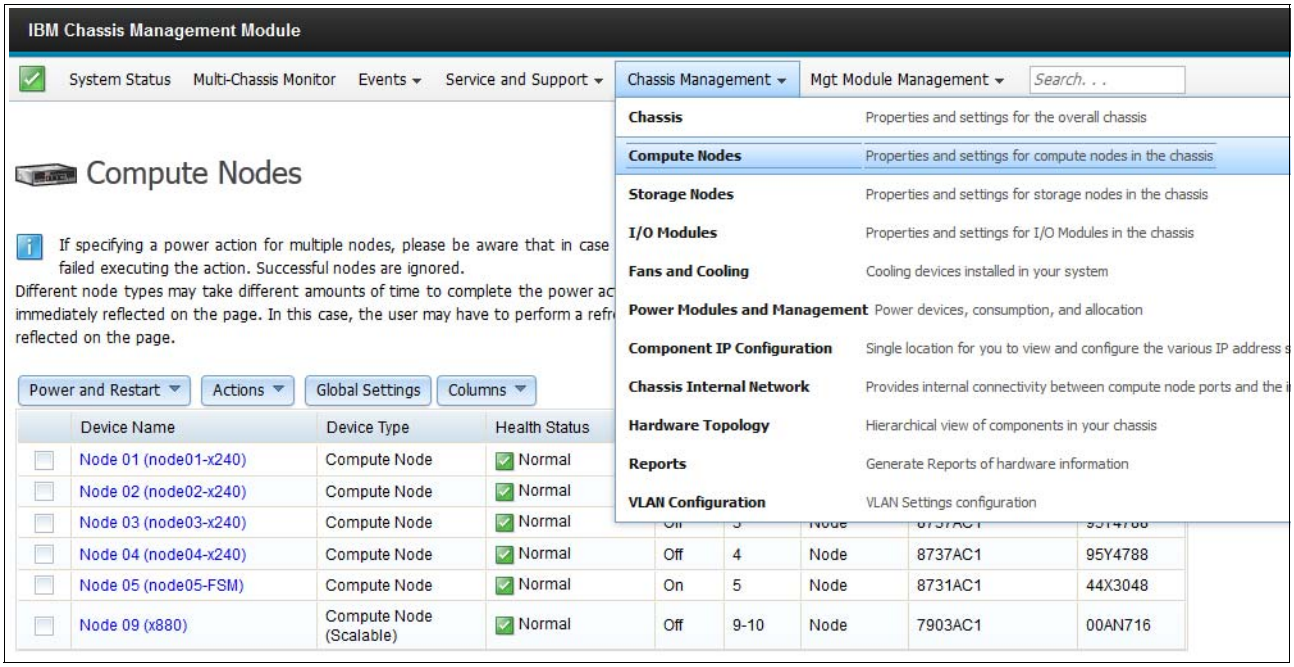

Figure 7-18 Chassis Management Compute Nodes menu

The Compute Nodes page also has a series of drop-down menus and buttons that feature the following functions:

- Power and Restart (node- or complex-specific):
  - Power On
  - Power Off
  - Shutdown OS and Power Off
  - Restart Immediately
  - Restart with Non-maskable Interrupt (NMI)
  - Restart System Management Processor
  - Boot to SMS Menu (Power Systems only)
- Actions (node- or complex-specific):
  - Launch Compute Node Console
  - Identify LED
- Settings (global across all installed nodes):
  - Policies:
    - Enable Local power control
    - Enable Wake on LAN
  - Serial over LAN: Enable Serial over LAN
- Columns (user interface display changes):
  - Device Name
  - Device Type
  - Health Status
  - Power
  - Bay

- Bay Type
- Machine Type Model
- I/O Compatibility
- WoL (Wake on LAN)
- Local Power Control
- Compute Expansion Module

Node- or complex-specific options require that a node is selected before the function can be applied.

Clicking one of the names in the Device Name column opens a window with details about that server, as shown in Figure 7-19.

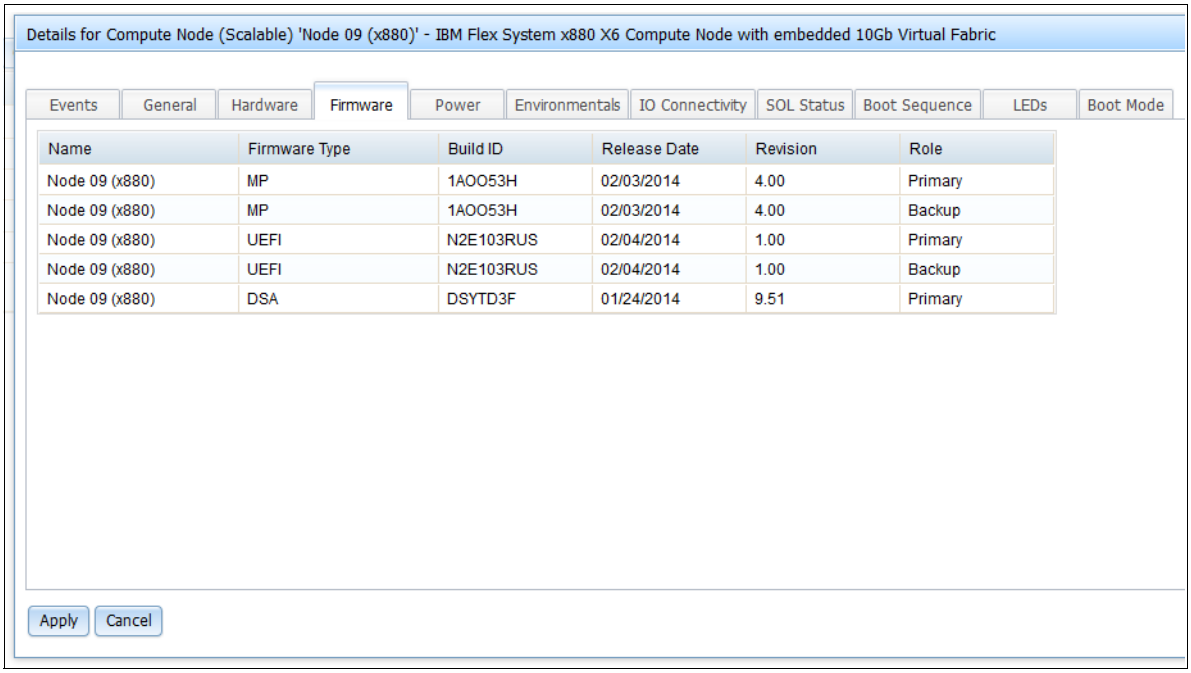

Figure 7-19 Compute Nodes details tab

7.5.4 Service and Support option

The Service and Support option is used for reviewing detected problems, troubleshooting, opening a service request, and for updating chassis settings.

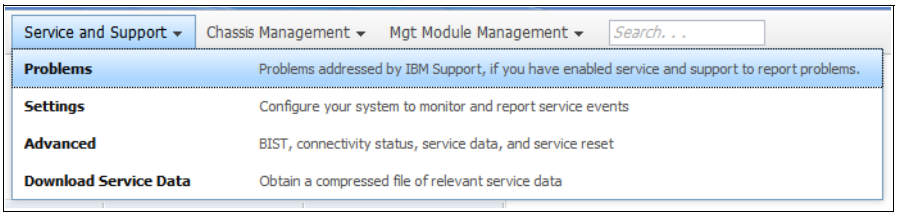

Figure 7-20 Service and Support tab

The Service and Support menu has four menu items:

Problems — Shows a grid of detected problems. You can open a Lenovo service request directly.

Settings	Use this menu option to configure the chassis, enter contact information, country, and settings, such as proxy access.
Advanced Status	This menu item provides advanced service information and more service tasks. You might be directed by Lenovo Support staff to review or perform tasks in this section.
Download Service Data	By using this menu item, you can download CMM data, send management module data to an email recipient (SMTP must be enabled), and download node data.

PureFlex System configurations: The Flex System Manager in a PureFlex System configuration can perform centralized reporting for all devices it manages, including the chassis components. Therefore, it is not necessary to configure this feature on the CMM.

Enabling Lenovo Support

The CMM call home feature is enabled and set up from the Settings option under the Service and Support menu bar option.

To enable Lenovo Support on the CMM, complete the following steps:

1. Click **Service and Support** → **Settings** from the menu bar option, as shown in Figure 7-21.

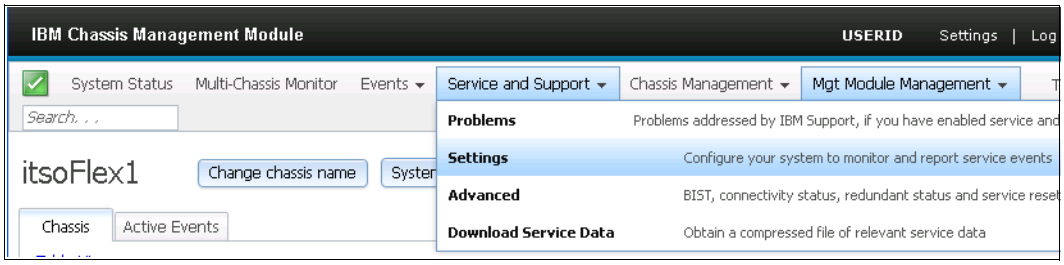

Figure 7-21 Selecting the CMM Service and Support Settings option

2. Read and acknowledge any licensing information that is presented to continue.

3. Complete the mandatory contact information, as shown in Figure 7-22.

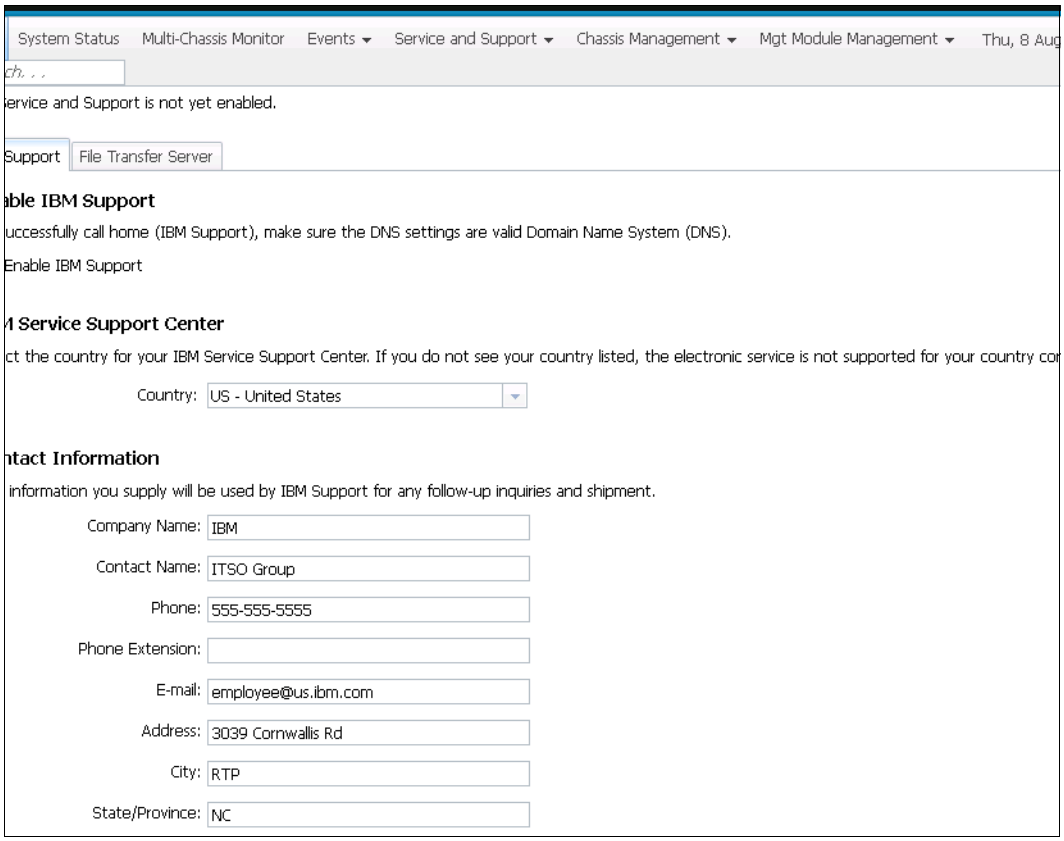

Figure 7-22 Required information to enable CMM phone home capability

4. Complete the optional information.
5. If a proxy is required for external communication to Lenovo Support, be sure to include this information in the optional settings, as shown in Figure 7-23.

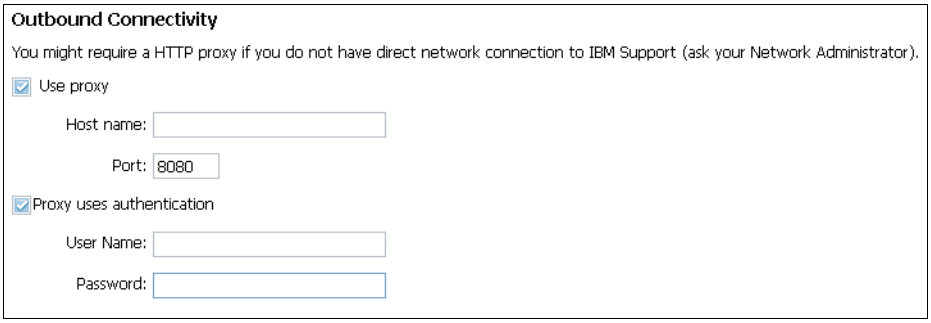

Figure 7-23 Optional information to enable CMM phone home capability

6. Click **Apply** to enable Lenovo Support and acknowledge any confirmation notices as they appear. Lenovo Support is now enabled, as shown in Figure 7-24.

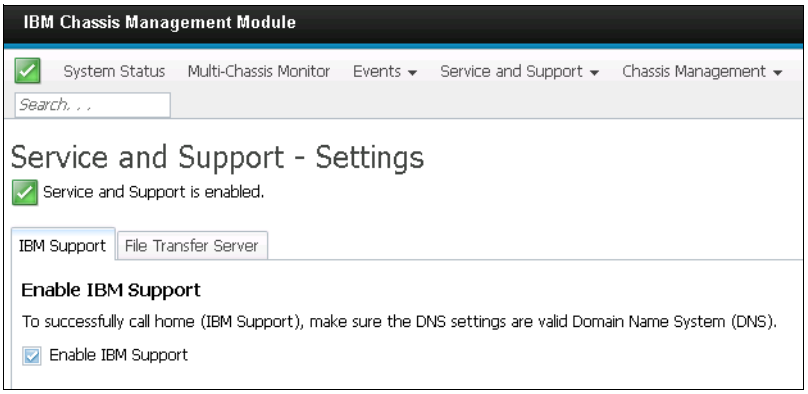

Figure 7-24 Lenovo Support enabled on CMM

7.6 Management by using Flex System Manager

This section describes the basic management of a X6 Compute Node by the Flex System Manager. The assumption is that the initial Flex System Manager setup wizard was run and at least one chassis with a X6 Compute Node was managed.

7.6.1 Accessing the Flex System Manager

Before you begin, you need the IP address of the Flex System Manager. You can access the Flex System Manager web interface by using a browser or the CLI from an SSH session. The browser method is described here.

For more information about supported browsers for accessing the Flex System Manager and all devices in the Flex System or PureFlex System, see “Planning for web browsers” in the PureFlex System information section on the Lenovo website:

`http://ibm.co/1sr8SLj`

Complete the following steps:

1. Open a browser and point it to the following URL (where `system_name` is the host name or IP address of the Flex System Manager):

   `https://system_name`

2. When the user login view displays, provide the proper user ID and password to complete the login process.
3. After login, the home tab view is displayed, as shown in Figure 7-25 on page 168.

All functions of the Flex System Manager can be accessed from this view with following second row of tabs:

- Initial Setup
- Additional Setup
- Plug-ins
- Administration
- Applications
- Learn

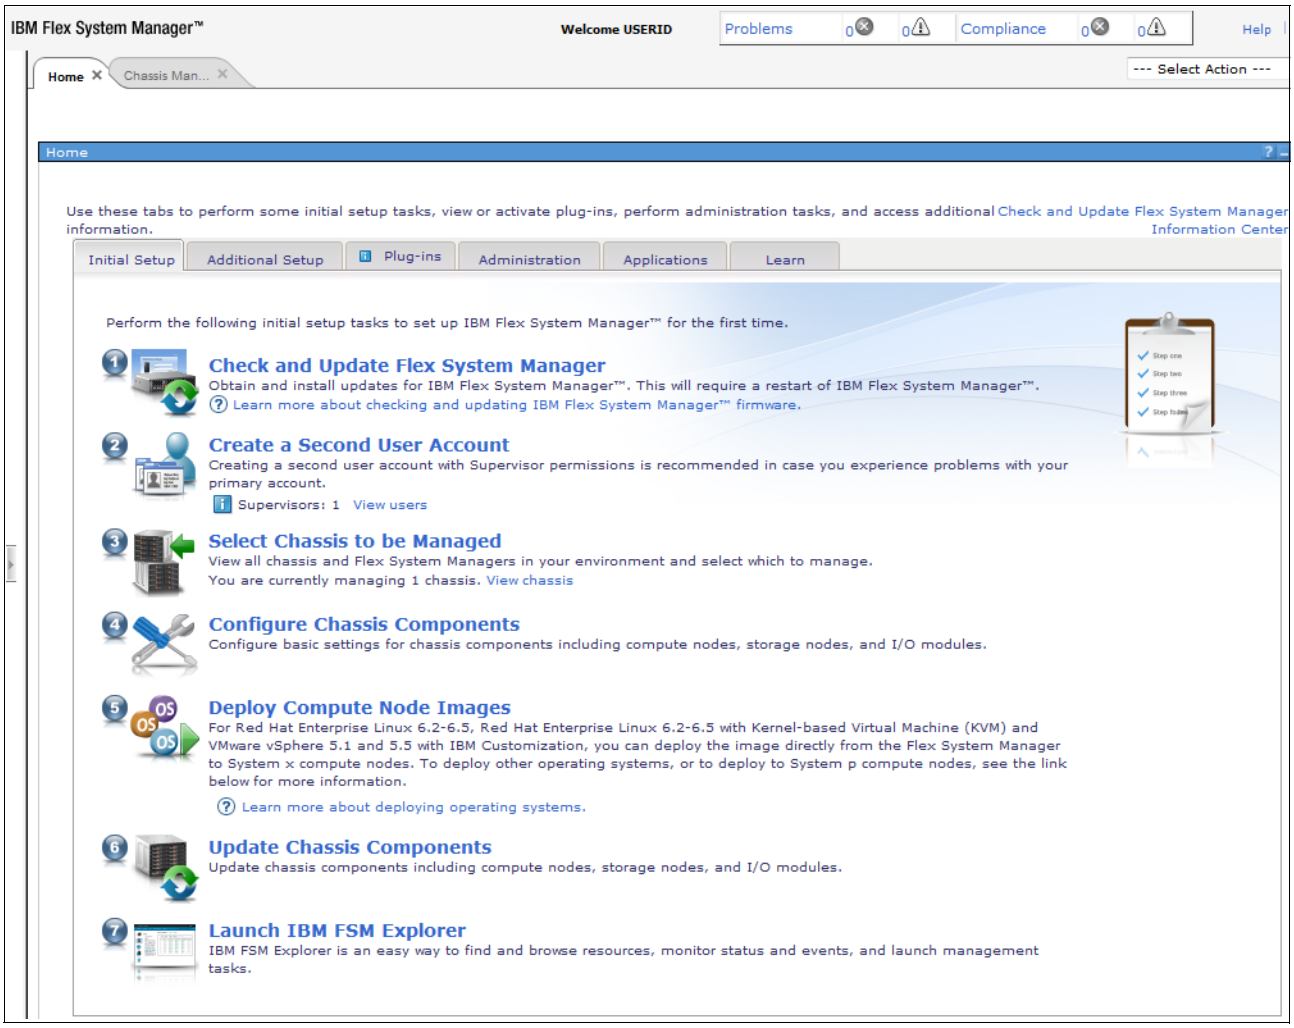

Figure 7-25 Flex System Manager Initial home page

7.6.2 Connecting a X6 Compute Node to the Flex System Manager

The following are dependencies when managing an X6 based compute node from the Flex System Manager:

- The CMM must successfully complete the discovery process of the node.
- The compute node's IP address is within the same subnet as the CMM.
- The Flex System Manager successfully managed the chassis containing the node.
- The Flex System Manager unlocked or successfully accessed the node's IMM.

The complete process for these dependencies is not described in this document but they are summarized next.

CMM discovery

When the chassis is powered up, the CMM restarted, or a compute node is inserted, a discovery process automatically occurs. This process establishes communications between the compute node and the CMM and allows the CMM to collect vital product data from the node.

During the chassis power up process or when a compute node is inserted, the power indicator light on the node fast flashes until the discovery process completes. When complete, the power indicator light is in a slow flash mode until power-on, and then it is on continuously.

The active chassis map that is shown on the CMM System Status status can also show the discovery mode when the mouse cursor is placed over the compute node image, as shown in Figure 7-12 on page 159.

Node IP configuration

The CMM Component IP Configuration option under Chassis Management is used to configure the IP addresses for the I/O modules, compute nodes, and storage nodes. These IP addresses are required to be in the same subnet as the CMM. For more information about how to configure a node, see “Component IP Configuration option” on page 160.

Flex System Manager chassis manage

After the Flex System Manager completes the initial configuration, the first task is to manage one or more chassis. This process establishes communication between the Flex System Manager and the target chassis CMM. During this process, the Flex System Manager authenticates with the CMM and collects initial chassis component vital product data. It also requests access (unlock) to the service processors in the various nodes and I/O modules in the chassis, including the IMM in X6 Compute Nodes.

Figure 7-26 shows the Flex System Manager Chassis Manager graphical view of a managed chassis with one x880 Compute Node in Bays 9 and 10.

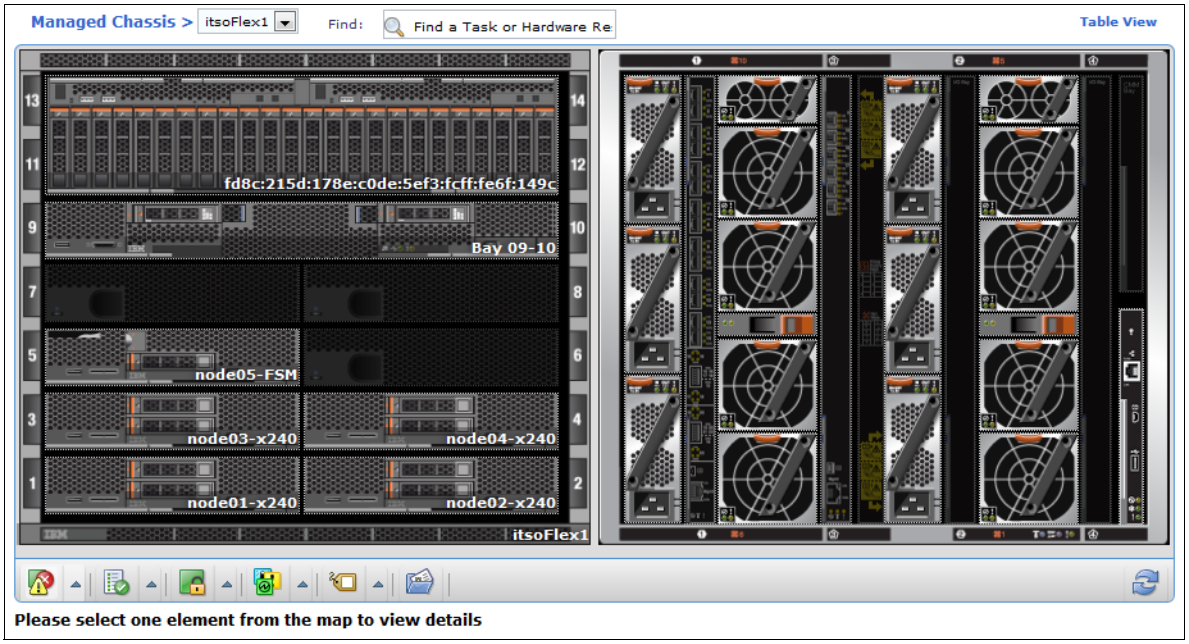

Figure 7-26 Flex System Manager discovered and managed Chassis graphical view

Flex System Manager compute node access

Figure 7-27 shows the same chassis in a table view. The table view has a column that is labeled “Access.” The wanted status is OK for compute and storage nodes and I/O modules. With this status, the Flex System Manager can communicate directly with the IMM in an X6 Compute Node.

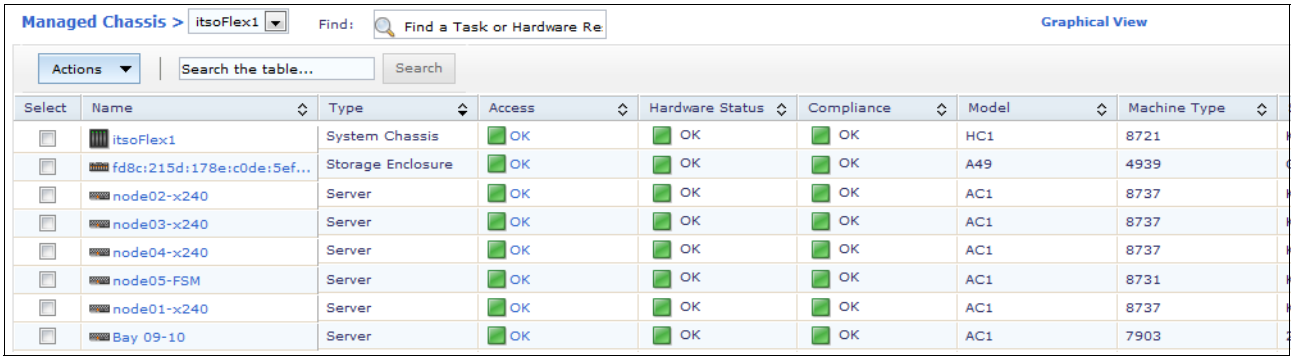

Figure 7-27 Flex System Manager discovered and managed chassis table view

Manage X6 resources navigation basics

The FSM Explorer view that is shown in Figure 7-28 is the starting point for basic X6 Compute Node management and can be reached by several methods, including the following most common methods:

- By clicking **Home** → **Plug-ins** → **Flex System Manager** → **Manage your Flex Resources**
- By clicking **Home** → **Initial Setup** tab → **Launch FSM Explorer**
- By clicking **Chassis Manager** → **General Actions** drop-down menu → **Resource Explorer**

This initial view shows the hardware or compute nodes that are currently known in all the managed chassis. This view has two areas of interest: a navigation list on the left side and the content area on the right side.

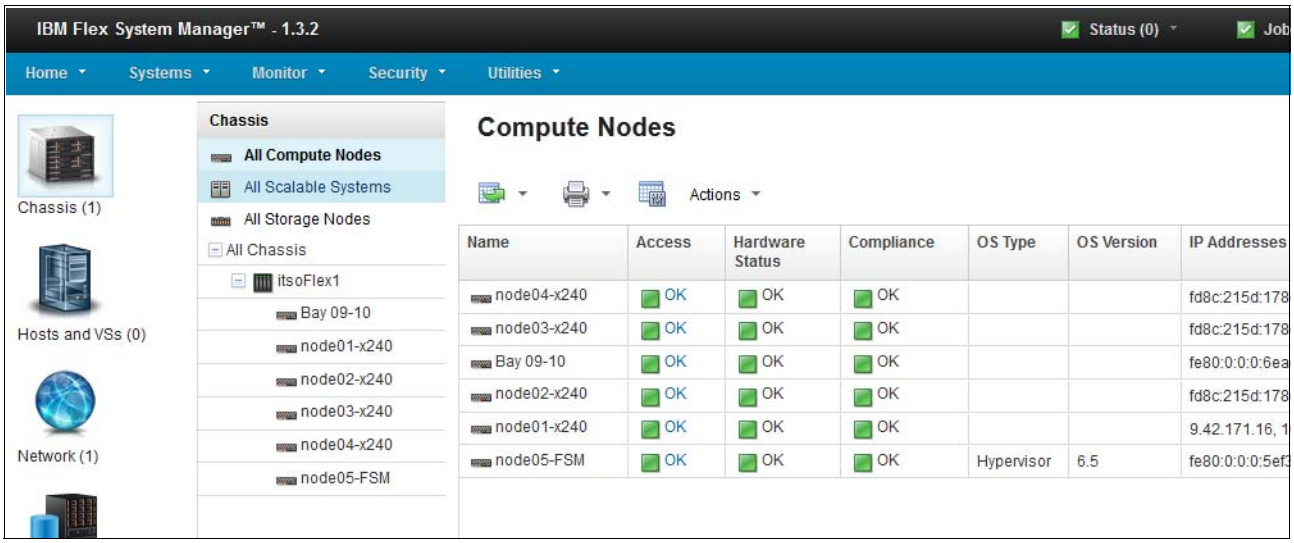

Figure 7-28 FSM Explorer initial screen

A list of scalable X6 Compute Nodes being managed by the Flex System Manager can be shown by clicking the **All Scalable Systems** option, as shown in Figure 7-29.

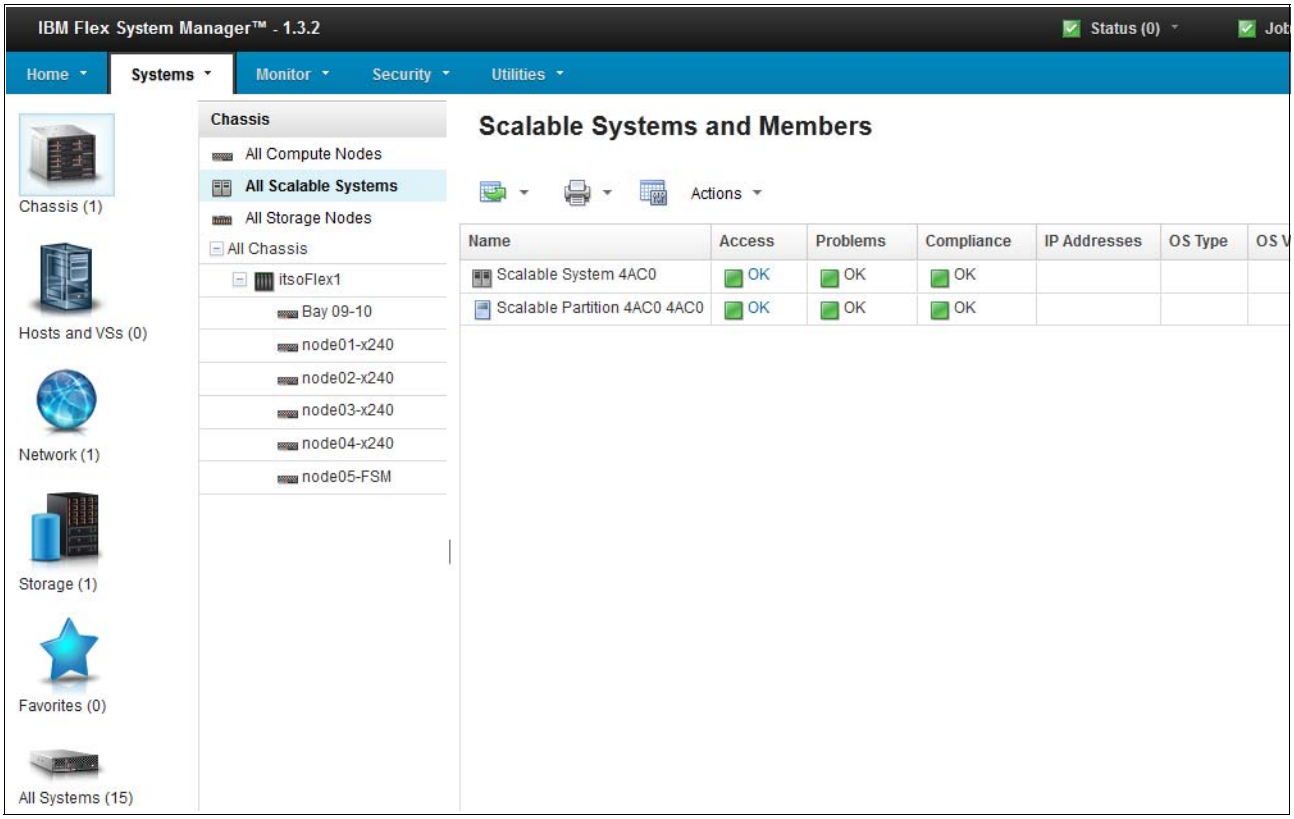

Figure 7-29 FSM Explorer Scalable Systems list

As shown in Figure 7-29, the left navigation options are used to directly access the following components:

- Chassis
  - Compute Nodes
  - Scalable Systems
  - Storage Nodes
- Hosts and Virtual Servers (VSs)
  - Virtual Servers
  - Hosts
- Network
  - Logical Networks
  - VLANs
  - Subnets
- Storage
  - Storage Pools
    - Storage Subsystems
    - Fabrics
- Favorites
- All Systems
  - Table view of all discovered logical and physical resources

For more information on Systems Management with the Flex System Manager, see the Lenovo Press publication, *Implementing Systems Management of IBM PureFlex System,* SG24-8060:

http://lenovopress.com/sg248060

7.7 Advanced Settings Utility (ASU)

The Advanced Settings Utility (ASU) allows you to modify your server firmware settings from a command-line interface environment. It supports multiple operating systems, such as Linux, Solaris and Windows (including Windows Preinstallation Environment). The firmware settings that can be modified on the X6 platform include UEFI and IMMv2 settings.

You can perform the following tasks by using the ASU:

- Modify the UEFI CMOS settings without the need to restart the system and access the F1 menu
- Modify the IMM setup settings
- Modify a limited set of vital product data settings
- Modify the iSCSI boot settings. To modify the iSCSI settings through ASU, you must initially manually configure the iSCSI settings through the server setup utility
- Remotely modify all of the settings through an Ethernet connection.

ASU supports scripting environments through batch-processing mode. Download the latest version and the *Advanced Settings Utility User's Guide* from the Advanced Settings Utility website:

http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-ASU

7.7.1 Using the ASU to configure settings in IMM

ASU 9.x supports configuring settings on servers with IMMs, such as the X6 Compute Nodes. The ASU uses the same set of commands and syntax that is used by previous versions of the ASU tool.

In IMM-based servers, you configure all firmware settings through the IMM. The ASU can connect remotely to the IMM over the LAN (out-of-band).

When the ASU runs any command on an IMM, it attempts to connect and automatically configure the LAN over a USB interface, if it detects that this interface is not configured. The ASU also provides a level of automatic and default settings. You have the option of specifying that the automatic configuration process is skipped, if you have manually configured, or do not want to have the IMM LAN over a USB interface.

See the *User's Guide for the Advanced Settings Utility* under Online Help on the following web page for more information:

http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-ASU

Tip: After you use the ASU to change settings, you must reset the IMM before you flash new firmware; otherwise, the changes to the settings might be lost. To reset the IMM, use the following ASU command:

```
asu rebootimm
```

Use the following procedure to download, install, and connect to the IMM2, using a Microsoft Windows operating system:

1. Create a directory named ASU.
2. Download the ASU tool for your operating system (32-bit or 64-bit) from the following web page, and save it in the ASU directory:

   http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-ASU
3. Unpack the utility:

   For Windows, double-click the file `ibm_utl_asu_asutXXX_windows_yyy.exe` file, which decompresses the files into the ASU folder.

   For Linux, open a terminal session and run the following command from the ASU directory:

   ```
   tar -zxf ibm_utl_asu_asutXXX_linux_yyyy.tgz
   ```
4. After it is unpacked, run a command, such as the **`asu show`** command, either in-band or out-of-band, using the commands that are listed in Table 7-2. This command confirms that the utility is installed and a connection has been established.

Table 7-2 Steps to run asu show both in-band and out-of-band methods

In-band	Out-of-band
Run the following command: ► For 32-bit: **`asu show`** ► For 64-bit: **`asu64 show`**	Ping the IMM to ensure that you have a network connection to the IMM. Run the following command: ► For 32-bit: **`asu show --host IPADDRESS`** ► For 64-bit: **`asu64 show --host IPADDRESS`**

7.7.2 Command examples

In this section, we provide a brief overview for the most commonly used ASU commands. See the *User's Guide for the Advanced Settings Utility* under Online Help on the following web page for all the ASU commands:

http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-ASU

Show all settings

On the command line, enter the **`asu show`** command. Example output is shown in Example 7-1.

Example 7-1 Output of the asu show command

```
IMM2.SSH_SERVER_KEY=Installed
IMM2.SSL_HTTPS_SERVER_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_HTTPS_SERVER_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_SERVER_DIRECTOR_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_SERVER_DIRECTOR_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_CLIENT_TRUSTED_CERT1=Not-Installed
IMM2.SSL_CLIENT_TRUSTED_CERT2=Not-Installed
IMM2.SSL_CLIENT_TRUSTED_CERT3=Not-Installed
IMM2.PowerRestorePolicy=Restore
IMM2.ThermalModePolicy=Normal
```

Show all UEFI settings

On the command line, enter the **asu show uefi** command. Example output is shown in Example 7-2.

Example 7-2 Output of the asu show uefi command

```
uEFI.OperatingMode=Custom Mode
uEFI.QuietBoot=Enable
uEFI.TurboModeEnable=Enable
uEFI.TurboBoost=Power Optimized
uEFI.ProcessorEistEnable=Enable
uEFI.ProcessorCcxEnable=Disable
uEFI.ProcessorC1eEnable=Enable
uEFI.HyperThreading=Enable
uEFI.EnableCoresInSbsp=All
uEFI.ExecuteDisableBit=Enable
uEFI.ProcessorVmxEnable=Enable
uEFI.ProcessorDataPrefetch=Enable
```

Show all IMM settings

On the command line, enter the **asu show imm** command. Example output is shown in Example 7-3.

Example 7-3 Output of the asu show imm command

```
IMM2.SSH_SERVER_KEY=Installed
IMM2.SSL_HTTPS_SERVER_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_HTTPS_SERVER_CSR=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CERT=Private Key and Cert/CSR not available.
IMM2.SSL_LDAP_CLIENT_CSR=Private Key and Cert/CSR not available.
```

Set TurboMode to enable

On the command line, enter **asu set uefi.turbomodeenable enable** command. Example output is shown in Example 7-4.

Example 7-4 Output of the ASU command setting turbomode to enabled

```
uEFI.TurboModeEnable=Enable
Waiting for command completion status.
Command completed successfully.
```

Disable Energy Manager

On the command line, enter **asu set uefi.energymanager disable** command. Example output is shown in Example 7-5.

Example 7-5 Output of the ASU command setting EnergyManager to disabled

```
uEFI.EnergyManager=Disable
Waiting for command completion status.
Command completed successfully.
```

7.8 MegaRAID Storage Manager

In this section, we provide an overview of the MegaRAID Storage Manager (MSM) software. For more information, see the *Installation and User's Guide* at the following web page:

http://ibm.com/support/entry/portal/docdisplay?lndocid=SERV-RAID

With MSM, you can configure, monitor, and maintain storage configurations on ServeRAID-M controllers. The MegaRAID Storage Manager GUI makes it easy for you to create and manage storage configurations. You can use MSM to manage local or remote RAID controllers and configure MSM for remote alert notifications. A command-line interface is also available.

To download the latest MegaRAID Storage Manager software and obtain the *Installation and User's Guide,* go to the ServeRAID software matrix web page:

http://ibm.com/support/entry/portal/docdisplay?lndocid=SERV-RAID

MegaRAID Storage Manager supports a variety of controllers. For a list of the most up-to-date information on operating system, server, and adapter support, see following web page:

http://www.ibm.com/support/entry/portal/docdisplay?Indocid=migr-5082327

7.8.1 MegaRAID Storage Manager installation

You must have administrator or root equivalent operating system privileges to install and fully access the MegaRAID Storage software.

There are four setup options:

- Complete: This option installs all program features.
- Client: This option installs the required components to remotely view and configure servers.
- Server: This option only installs the required components for remote server management.
- Stand-alone: This option only installs the required components for local server management.

7.8.2 Drive states

A *drive group* is one or more drives that are controlled by the RAID controller.

There are multiple drive states. The following list describes all of the possible drive states:

Online	A drive that can be accessed by the RAID controller and is part of the virtual drive.
Unconfigured Good	A drive that is functioning normally but is not configured.
Hot Spare	A drive that is powered up and ready for use as a spare in case an online drive fails. There are two Hot Spare drive states: Dedicated and Global.
Failed	A drive that was originally configured as Online or Hot Spare but one on which the firmware detects an unrecoverable error.
Rebuild	A drive to which data is being written to restore full redundancy for a virtual drive.

Unconfigured Bad	A drive on which the firmware detects an unrecoverable error.
Missing	A drive that was Online but which has been removed from its location.
Offline	A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned.

7.8.3 Virtual drive states

A *virtual drive* is a partition in a drive group that is made up of contiguous data segments on the drives.

There are multiple virtual drive states. The following list describes all of the possible virtual drive states:

Optimal	The virtual drive operating condition is good. All configured drives are online.
Degraded	The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
Partially Degraded	The operating condition in a RAID-6 virtual drive is not optimal. One of the configured drives has failed or is offline.
Failed	The virtual drive has failed.
Offline	The virtual drive is not available to the RAID controller.

7.8.4 MegaCLI utility for storage management

In this section, we provide an overview of the MegaCLI utility. For more information, see the *Installation and User's Guide* at the following web page:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5073015

The MegaCLI utility is a CLI application. You can use this utility to configure, monitor, and maintain ServeRAID SAS RAID controllers and the devices that connect to them.

Creating a virtual drive with command-line interface

In this example, we have two hard disk drives in slots one and two. Both hard disks must be in *Unconfigured Good* state before configuration of a virtual drive.

Follow these steps to create a virtual drive by using the CLI:

1. Use the following command to locate the Enclosure Device ID and the Slot Number of both hard disk drives:

   ```
   MegaCli -PDList -aAll
   ```

2. Now we can create the virtual drive. In our example, we issue the following command to create a RAID-1:

   ```
   MegaCli -CfgLDAdd -R1[252:1,252:2] -a0
   ```

   Example 7-6 shows the resulting output.

Example 7-6 Output from command MegaCli -CfgLDAdd -R1[252:1,252:2] -a0

```
Adapter 0: Created VD 1

Adapter 0: Configured the Adapter!!
```

```
Exit Code: 0x00
```

3. The virtual drive is successfully created.

Additional command examples

The following command examples use the MegaCli command:

- Display help for MegaCLI

  ```
  MegaCli –h|–Help|?
  ```

- Display controller properties for all installed adapters

  ```
  MegaCli -AdpAllinfo -aALL
  ```

- Save configuration on the controller

  ```
  MegaCli -CfgSave -f c:\saveconfig.txt -a0
  ```

- Restore configuration data from file

  ```
  MegaCli -CfgRestore -f c:\saveconfig.txt -a0
  ```

- Display virtual drive information for all VD on all adapters

  ```
  MegaCli -LDInfo -Lall -aALL
  ```

- Display virtual drive and physical drive information for all adapters

  ```
  MegaCli -LDPDInfo -aAll
  ```

- Display number of virtual drives for all adapters

  ```
  MegaCli –LDGetNum -aALL
  ```

- Display list of physical devices for all adapters

  ```
  MegaCli –PDList -aAl
  ```

Abbreviations and acronyms

AC	alternating current
ACID	atomicity, consistency, isolation, and durability
ACPI	advanced control and power interface
ADJ	adjacent
AES	Advanced Encryption Standard
AES-NI	Advanced Encryption Standard New Instructions
APIC	Advanced Programmable Interrupt Controller
ASHRAE	American Society of Heating, Refrigerating and Air Conditioning Engineers
ASIC	application-specific integrated circuit
ASU	Advanced Settings Utility
AVX	Advanced Vector Extensions
BIOS	basic input/output system
BM	Bridge module
BoMC	Bootable Media Creator
BRD	board
CD	compact disk
CFM	cubic feet per minute
CIM	Common Information Model
CLI	command-line interface
CMM	Chassis Management Module
CMOS	complementary metal oxide semiconductor
CNA	Converged Network Adapter
COD	configure on disk
CPU	central processing unit
CRC	cyclic redundancy check
CRM	Customer Relationship Management
CTO	configure-to-order
DC	domain controller
DCU	data cache unit
DDDC	Double Device Data Correction
DDF	Disk Data Format
DDR	Double Data Rate
DHCP	Dynamic Host Configuration Protocol
DIMM	dual inline memory module
DRAM	dynamic random access memory
DSA	Dynamic System Analysis
DVD	Digital Video Disc
DWPD	drive writes per day
ECC	error checking and correcting
ERP	enterprise resource planning
FC	Fibre Channel
FDR	fourteen data rate
FoD	Features on Demand
FSM	Flex System Manager
FSP	flexible service processor
FTP	File Transfer Protocol
GB	gigabyte
GT	Gigatransfers
GUI	graphical user interface
HBA	host bus adapter
HDD	hard disk drive
HPC	high-performance computing
HS	hot swap
HTTP, http	Hypertext Transfer Protocol
HVDC	high voltage direct current
I/O	input/output
IBM	International Business Machines
ID	identifier
IMM	Integrated Management Module
IOPS	I/O operations per second
IP	Internet Protocol
IPMI	Intelligent Platform Management Interface
IRQ	interrupt request
ISO	International Organization for Standards
IT	information technology
ITSO	International Technical Support Organization
JBOD	just a bunch of disks
KVM	kernel-based virtual machine (or) keyboard-video-mouse
LAN	local area network

LDAP	Lightweight Directory Access Protocol
LED	light emitting diode
LOM	LAN on Motherboard
LP	low profile
LR-DIMM	load-reduced DIMM
LUN	logical unit number
MAC	media access control
MB	megabyte
MCA	Machine Check Architecture
MEM	memory
MIS	management information system
MLC	multi-level cell
MSI	Microsoft Installer
MSM	MegaRAID Storage Manager
MT/s	millions of transfers per second
NFS	network file system
NIC	network interface card
NL	nearline
NMI	non-maskable interrupt
NOS	network operating system
NPIV	N_Port ID Virtualization
NUMA	Non-Uniform Memory Access
NVRAM	non-volatile random access memory
OEM	other equipment manufacturer
OLAP	online analytical processing
OLTP	online transaction processing
OVF	Open Virtualization Format
PCH	Platform Controller Hub
PCI	Peripheral Component Interconnect
PE	Preinstallation Environment
PFA	Predictive Failure Analysis
POST	power-on self-test
PXE	Preboot eXecution Environment
QPI	QuickPath Interconnect
RAID	redundant array of independent disks
RAM	random access memory
RAS	remote access services; row address strobe; reliability, availability, and serviceability
RBS	redundant bit steering
RDIMM	registered DIMM
RDMA	Remote Direct Memory Access
RHEL	Red Hat Enterprise Linux
ROC	RAID-on-card, RAID-on-Chip
ROM	read-only memory
RSS	Receive-side scaling
SAN	storage area network
SAS	Serial Attached SCSI
SATA	Serial ATA
SCM	Supply Chain Management
SDDC	Single Device Data Correction
SED	self-encrypting drive
SFF	Small Form Factor
SLES	SUSE Linux Enterprise Server
SMI	scalable memory interconnect
SMS	System Management Services
SMTP	simple mail transfer protocol
SMX	Safer Mode Extensions
SNMP	Simple Network Management Protocol
SQL	Structured Query Language
SRIOV	single root I/O virtualization
SS	simple swap
SSD	solid-state drive
SSH	Secure Shell
SSL	Secure Sockets Layer
TB	terabyte
TCP/IP	Transmission Control Protocol/Internet Protocol
TDP	thermal design power
TFTP	Trivial File Transfer Protocol
TOE	TCP offload engine
TTY	teletypewriter
TXT	text
UEFI	Unified Extensible Firmware Interface
URL	Uniform Resource Locator
USB	universal serial bus
UXSP	UpdateXpress System Pack
UXSPI	UpdateXpress System Pack Installer
VAC	Volts alternating current
VD	virtual drive
VFA	Virtual Fabric Adapter
VLAN	virtual LAN
VM	virtual machine
VT	Virtualization Technology

Related publications

The publications listed in this section are particularly suitable for a more detailed discussion of the topics covered in this book.

Lenovo Press publications

The following Lenovo Press publications provide additional information about the topic in this document.

- *Flex System X6 Compute Nodes* Product Guide
  http://lenovopress.com/tips1160
- *Flex System Products and Technology*, SG24-8255
  http://lenovopress.com/sg248255
- *Flex System Interoperability Guide*:
  http://lenovopress.com/fsig
- *Flex System Networking in an Enterprise Data Center*, REDP-4834
  http://lenovopress.com/redp4834
- *Migrating from BladeCenter to Flex System*, REDP-4887
  http://lenovopress.com/redp4887
- *System x3850 X6 and x3950 X6 Planning and Implementation Guide*, SG24-8208
  http://lenovopress.com/redp8208
- *xREF: System x Reference*:
  http://lenovopress.com/xref

Online resources

The following websites are also helpful for further information:

- Flex System documentation:
  http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp
- *Flex System Enterprise Chassis Power Requirements Guide*:
  http://ibm.com/systems/bladecenter/resources/powerconfig.html
- *Configuration and Option Guide*:
  http://www.ibm.com/systems/xbc/cog/
- ServerProven
  http://ibm.com/systems/info/x86servers/serverproven/compat/us/flexsystems.html
- ServerProven operating system support for Flex System
  http://ibm.com/systems/info/x86servers/serverproven/compat/us/nos/flexmatrix.shtml

- IBM System Storage Interoperation Center:
  http://www.ibm.com/systems/support/storage/ssic
- *Integrated Management Module II User's Guide*:
  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086346

Lenovo Flex System X6 Compute Node (7903) Planning and

Learn about the Flex System x880 X6, x480 X6 and x280 X6 Compute Nodes

Read our guidance about scalability and paritioning

Be informed about our latest options and features

The increasing demand for cloud computing and business analytic workloads drives innovation to find new ways to build information systems. Lenovo clients are looking for cost-effective IT solutions that manage large amounts of data, easily scale performance, and provide reliable real-time access to actionable information.

Built on decades of innovation, Lenovo introduces its sixth generation of X-Architecture technology, the X6 servers, which includes the Flex System x880 X6, Flex System x480 X6, and Flex System x280 X6. These X6 servers are fast, agile, and resilient:

- Fast application performance means immediate access to actionable information.
- Agile system design helps reduce acquisition costs and provides the ability to host multiple generations of technology in a single server.
- Resilient platforms maximize application uptime and offer easy integration into virtual environments.

This Lenovo Press book describes the product and covers planning and implementation. The first chapters provide technical information about the 2-socket x280 X6, 4-socket x480 X6, and 8-socket x880 X6. This information is most useful in designing, configuring, and planning your server order. The later chapters provide configuration and setup details to get your server operational.

This book is for Lenovo clients, Business Partners, and employees who want to understand the features and capabilities of the X6 portfolio of servers and want to learn how to install and configure the servers for use in production. It is related to the Lenovo Press book about the rack-mount X6 servers, *System x3850 X6 and x3950 X6 Planning and Implementation Guide*, SG24-8208.

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

Lenovo

Lenovo Flex System X6 Compute Node (7903) Planning and Implementation Guide

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages